STANFORD UNIVERSITY

CENTER FOR SYSTEMS RESEARCH

BIOCYBERNETIC FACTORS IN HUMAN PERCEPTION AND MEMORY

James Anliker
Robert Floyd
Martin Morf
Thomas Kailath

Information Systems Laboratory

FINAL REPORT

for the Advanced Research Projects Agency
of the Department of Defense

BIOCYBERNETIC FACTORS IN HUMAN PERCEPTION AND MEMORY

James Anliker
Robert Floyd
Martin Morf
Thomas Kailath

ISL-BIOCYB-77

Contract Number: N00014-76-C-0597
ARPA Order Number: 3177

July 1977

134 p.

Department of Electrical Engineering
Stanford Electronics Laboratories
Stanford University Stanford, California 94305

UNCLASSIFIED

RITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
RT NUMBER ISL-BIOCYB-77	2. GOVT ACCESS ON NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) BIOCYBERNETIC FACTORS IN HUMAN PERCEPTION AND MEMORY		5. TYPE OF REPORT & PERIOD COVERED Final, 2-1-76 - 4-30-77 6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) James Anliker, Robert Floyd, Martin Morf, Thomas Kailath		8. CONTRACT OR GRANT NUMBER(s) N00014-76-C-0597
9. PERFORMING ORGANIZATION NAME AND ADDRESS Stanford University Dept. of Electrical Engineering Stanford, CA. 94305		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS NR201-229 000000
11. CONTROLLING OFFICE NAME AND ADDRESS Procuring Contracting Officer Office of Naval Research, Dept. of the Navy 800 N. Quincy St., Arlington, Virginia 22217		12. REPORT DATE July 1977 13. NUMBER OF PAGES 117
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office) Same as 11		15. SECURITY CLASS. (of this report) UNCLASSIFIED 15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)
Approved for public release; distribution unlimited

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)
Approved for public release; distribution unlimited

18. SUPPLEMENTARY NOTES
Final Report

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)
Biocybernetics, eye-movements, brain waves, eye-tracking, man-machine systems, memory, alpha phase, fast algorithms.

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)
This project was concerned with the development of biocybernetic concepts and methods which have potential value for enhancing visual perception of, and visual memory for, scenic materials. The essence of the biocybernetic idea is that feedback and control schemes implemented by various machines can be used to enhance or extend various aspects of human performance beyond unaided limits. To design effective closely-coupled man-machine systems it is desirable to have the perceptual and behavioral aspects of human performance formulated in terms of feedback and control principles. In this project the investigators concentrated

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102 LF 014 6601

UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

on studying the feedback and control possibilities inherent in the coupling of visual stimuli to eye-direction and to the phase of the EEG alpha rhythm. They succeeded in developing state-of-the-art systems for real-time tracking of eye-direction and alpha-phase; these computerized tracking systems are capable of controlling visual stimuli so that their occurrence is conditional upon eye-direction and alpha-phase.

Their computerized eye-tracking system, known as PERSEUS, incorporates the 2-dimensional double-Purkinje-image eye-tracker (DPIET) developed at the Stanford Research Institute. This non-contacting device tracks and compares the positions of the first and fourth Purkinje images formed by reflections of infrared light beamed at the subject's eye. By comparing the position of the fourth Purkinje image with respect to the first Purkinje image, one obtains a sensitive measure of eye-rotation which is uncontaminated by eye-translation. PERSEUS, which is implemented in a PDP-15/76 computer, provides computer-generated scenic targets for experimental studies, calibration routines for the linearization of eye-direction estimates obtained from the DPIET, and real-time analysis of fixation and saccades in relation to scenic targets. PERSEUS is capable of stabilizing a computer-generated scenic image upon the retina of the moving eye; it is also capable of modifying the properties computer-generated scenic target on the basis of current eye-direction. A model for saccadic eye-movement was developed which permits real-time prediction of saccadic destination from the early portion of the saccadic trajectory.

A computerized scheme was implemented for the phase-contingent analysis of the average visual evoked potential. The results of this investigation were found to be compatible with the hypothesis that cerebral incorporation of visual sensory inputs is based upon phase-modulation of the EEG alpha rhythm. A more elaborate analysis of the average visual evoked potential in terms of alpha phase probability at the moment of photic stimulation also produced results consistent with the phase-modulation hypothesis. According to the phase-modulation theory advanced by the investigators, the cerebral encoding of visual sensory inputs is based on the phase-difference between the cortical alpha phase based upon autonomous stimulation by the thamalic pacemaker and the observed cortical phase which may be shifted by the action of the sensory (exogenous) input. According to this theory, the alpha 'carrier' would have to be activated in order for demodulation of mnemonic playback to occur. A nonlinear model of the human EEG signal was developed and tested.

Since practical success in achieving a biocybernetic 'close coupling' of man and machine depends upon an enlighted selection of suitably efficient modeling techniques, a review of recursive modeling methods was undertaken to provide a systematic classification of an area characterized by rapid and diverse developments.

FOREWORD

The Stanford Biocybernetics Project was begun in 1972 with Professor D. C. Lai and Professor T. Kailath as co-principal investigators. Dr. J. Anliker of the NASA/Ames Research Center played an active role as consultant in the planning of this project, made the facilities of his laboratory available to the project, and helped to guide psycho-physiological and experimental aspects of the project. Dr. H. Magnuski served as Post-Doctoral Research Fellow during 1973-1974, while A. Huang, K. Jacker and L. Strickland provided programming assistance. Two Ph.D. theses by J. Nickolls and Arun Shah, both Graduate Student Research Assistants, were completed under the project. After Professor Lai's return to Vermont in 1975, the major effort was contributed by Dr. Anliker and by Mr. R. V. Floyd who served as Scientific Programmer. Professor Martin Morf made valuable contributions in the development of fast algorithms.

T. Kailath

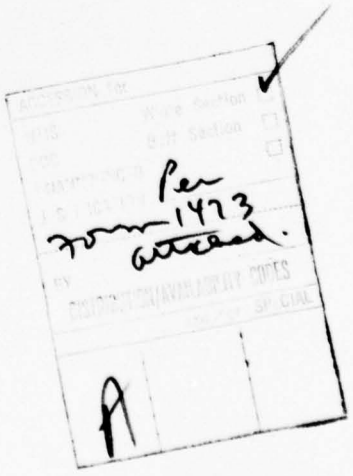

PROFESSIONAL PERSONNEL

1. Dr. Thomas Kailath, Professor of Electrical Engineering
   and Director of the Information Systems Laboratory
   Principal Investigator

2. Dr. Martin Morf, Assistant Professor of Electrical Engineering

3. Dr. James Anliker, Visiting Scholar in Electrical Engineering

4. Robert V. Floyd, Graduate Student in Neurosciences Program
   Senior Computer Programmer

5. John R. Nickolls, Graduate Student in Electrical Engineering

6. Arun Shah, Graduate Student in Electrical Engineering

ABSTRACT

The ARPA-sponsored Biocybernetics Project at Stanford University has been concerned with the development of biocybernetic methods for the enhancement of visual perception of, and memory for, scenic materials. The essence of the biocybernetic concept is that feedback and control schemes implemented by appropriate machines can be used to enhance or extend various aspects of human performance beyond the unaided limits. In the present project the focus has been upon enhancing human perception and memory--in particular, perception and memory of scenic stimuli. Since it is recognized that some individuals are decidedly better visualizers than others and since it is thought that the ability to visualize scenic materials would confer definite advantages, this project was conceived to explore the possibility that biocybernetic methods might be employed to good effect in the study of human visualization processes and to determine whether such skills might be trainable through appropriate feedback or extended by computerized prosthetic devices.

In this project we have concentrated on analyzing the feedback and control possibilities inherent in eye-pointing behavior and in the electroencephalographic alpha rhythm. To this end we have succeeded in developing what is currently the most advanced eye-tracking system in existence, and, in

collaboration with the Ames Research Center, we have also developed a state-of-the-art computerized system for tracking the EEG alpha rhythm. Both of these systems have the capability of monitoring their biological targets on an on-line, real time basis. In addition, they can control visual stimulation so that its occurrence is conditional upon current states in the biological targets. These systems are described in some detail in the body of this report, its appendices, and references. Work on this project has also been reported in various publications, reports, and theses; two Ph.D. dissertations have been submitted to the Department of Electrical Engineering, and a third Ph.D. dissertation is being prepared by a student in the Neurosciences Program.

Accurate real-time monitoring of human eye-pointing constitutes a very considerable technical task. Our advanced computerized eye-tracking system, known as PERSEUS, incorporates the 2-dimensional double-Purkinje-image eye-tracker (DPIET) described by Cornsweet and Crane (1973) and updated by Crane and Steele (unpublished report, Stanford Research Institute). The DPIET is a hardware system consisting of an infrared light source, servo-controlled optics, and associated electronic circuitry. This device tracks and compares the positions of the first and fourth Purkinje images formed by reflections of infrared light beamed at the subject's eye. Inasmuch as these two Purkinje images change their separation when the eye rotates but move the same amount

in the same direction when the eye translates, their amount of separation provides a sensitive measure of eye-rotation which is uncontaminated by the principal source of eye-tracking error, viz., eye-translation. However, the DPIET has certain limitations which must be compensated for by the computerized system. For instance, the DPIET knows nothing about higher order phenomena such as fixations and saccades; these must be discriminated as temporal patterns appearing in the continuous output voltages representing horizontal and vertical eye-positions. Nor does the DPIET have any provision for creating and controlling the display of visual stimuli. Finally, there are the idiosyncrasies of each subject's eye which can produce nonlinear distortions in the eye-tracker estimates of eye-pointing coordinates. In PERSEUS the analysis of fixations and saccades, the generation and control of scenic stimuli, and the linearization of each subject's nonlinearities are all handled by a medium-sized disk-oriented digital computer, the PDP-15/76. This highly accurate (less than 4 minutes of arc error over a field 20 degrees in diameter) eye-tracking system has been used to implement various biocybernetic schemes for controlling scenic displays on the basis of eye-position. It is capable, for example, of stabilizing a scenic stimulus upon the subject's retina without any attachments to the eye. While in this mode it can blank any portion of the display, e.g., periphery, parafovea, or fovea. PERSEUS performs real-time analyses on sequential fixations and saccades, numbers scanpath fixations, and superimposes this information on an

ancillary CRT display of the visual stimulus. The data collected during experiments are stored and subjected to more complex forms of off-line analysis. Of particular interest are the visual contact probability plots; these 3-dimensional surfaces are formed by weighting the area surrounding each fixation by an estimate of the local retinal acuity. For quantitative purposes the contact probability surfaces are sliced at equal amplitude intervals to produce a set of iso-contours which are then superimposed upon the scenic stimulus pattern. By comparing these iso-contours with a similar analysis of the spatial frequencies in the stimulus it is possible to obtain a measure of selective attention which takes into account the eye's natural preference for high frequency "hot spots."

Our studies of the EEG alpha rhythm have been of two kinds: (a) detailed off-line analyses of the individual sample functions which are ordinarily subjected to time averaging to produce the visual average evoked potential and (b) studies of phase-conditional photic stimulation. For the evoked potential studies we have employed quadrature analysis for the definition of EEG alpha phase, i.e., phase with respect to a coherent alpha phase (<u>not</u> interhemispheric phase differences). In one study we examined the role of the contingent alpha phase, i.e., the phase of the alpha rhythm at the moment of photic stimulation. When we reclassified the AEP sample functions into twenty different contingent phase ranges, averaged the sample functions within

each phase range separately, and compared these contingent averages, we found that the stimulus seemed to phase-shift each of these averages in the latency range associated with the prominent N1-P2-N2 complex in the AEP. This suggested that the stimulus might be acting to phase-shift the cortical alpha rhythm away from the phase anticipated in the autonomous or unstimulated alpha rhythm. This would produce a cortical phase at variance with the thalamic pacemaker. To check further into this possibility, we extracted the time-varying phase function from each of the raw EEG sample functions. This gave us the phase of the cortical alpha activity in each sample function at 2-msec intervals following the stimulus. From these phase data we constructed 3-dimensional phase probability surfaces (phase, latency, probability) for each of the phase-contingent subsets of the AEP data. Then we extracted the modal phase function from each of these contingent phase probability surfaces. This analysis showed clearly that our previous conclusion based on phase contingent averaging was correct. That is, the photic stimulus acts shift the contingent phases in the direction of a "preferred" phase at the latency corresponding to the prominent N1-P2-N2 complex in the AEP. These results are also consistent with our phase uncertainty model of stimulus encoding in the visual cortex.

For our phase-conditional stimulation studies we devised a novel method for tracking and modeling the frequency and phase of the cortical alpha rhythm. In this computerized

scheme, the alpha frequency is monitored continuously by autocorrelating the EEG signal and interpreting this transform. This computation is accomplished by a special purpose computer (SAICOR 400-pt. correlator). A general purpose computer does the interpretation into frequency and uses this information to control the frequency parameter of a programmable oscillator. The output of the oscillator (now tuned to the alpha frequency) is fed into another correlator where it is cross correlated with the current EEG signal. The output of this correlator is interpreted as phase difference between the two signals; this phase information is used to null the phase difference by adjusting the phase parameter of the programmable oscillator. Through this complex procedure the output of the programmable oscillator is made to model closely the frequency and phase properties of the EEG alpha rhythm. From this point it is a relatively simple matter for the computer to make the presentation of a photic stimulus conditional upon a specified phase of the alpha rhythm.

In a project such as this, closely-coupled systems cannot be achieved without the aid of efficient estimation and prediction of the brain's states as revealed in such outputs as eye-movements and EEG scalp potentials. Although the brain's activities are complex and time-varying, linear systems modeling techniques can be used successfully to predict brain states. Model parameter estimates can be updated in time, yielding a useful time-varying linear model.

Nonstationarities in observed outputs can be modeled with time-varying parameters, or with a time-invariant linear model with suitable initial conditions. In the appendix entitled "A Classification of Recursive Modeling Methods" we have presented a systematic classification of exact least-squares modeling procedures that are recursive (in model-order) and optimal in some sense. Methods which are sub-optimal or approximate are only briefly indicated.

Of the three types of procedures considered, e.g., Riccati or square-root methods, fast methods utilizing matrix shift-invariance properties, and ladder-form methods, the latter two are most suited to problems requiring efficiency. In particular, the recursive (in time) versions lend themselves to on-line or real-time applications because the input-output data are assessed sequentially one sample at a time. Also, new parameter estimates are available at each sample time, which facilitates the solution of on-line control problems.

The ladder-form methods also have the desirable property that their stability can be checked merely by inspection of the reflection coefficients. In addition, they are numerically robust since the major operations are sample cross correlations. They also have minimal storage requirements for least-squares modeling methods. The structure of the ladder-type realization suggests that the reflection coefficients may have physical significance in the process being modeled.

Finally, it may be noted that practical success in achieving the desired close-coupling of man and machine will depend upon an enlightened selection of suitably efficient modeling techniques. Our review of recursive modeling methods is intended to provide a systematic classification of an area characterized by rapid and diverse developments.

INTRODUCTION

The purpose of the Stanford Biocybernetics Project was to investigate the feasibility of developing and using biocybernetic technology (a) to study visual picture perception and scenic memory and (b) to enhance human memory--in particular the quality, persistence, and retrievability of visual images of environmental objects. Would it be possible, in other words, to use real-time computerized monitors and feedback loops to arrange extraordinary coincidences of eye-movement *scanpaths*, electroencephalographic states, and visual stimulus configurations so as to favor the production of unusually vivid, persistent, and memorable visual images of the stimulus material? Although the common adult experience with post-stimulus visual images is one of rapid fading, there have been persistent reports that, in some individuals, vivid visual images persist much longer than is usual and that they can, in these cases, sometimes be recalled with remarkable accuracy and modified at will. Some investigators have viewed these phenomena as evidence of neural pathology while others have tended to regard them as manifestations of unique biological gifts. We are inclined to ask whether they may be regarded as indications that these individuals simply utilize more successful strategies for the impression, storage, and retrieval of scenic (as opposed to verbal) information.

Available evidence suggests that unusually strong visual imagery is most frequently observed in children and only rarely persists into adulthood; furthermore, it appears that the development and exercise of verbal skills coincides with a reduced incidence of exceptional visual imagery of the eidetic type. If it could be shown that the basic capacity for visual memory is a common human property which is given a low priority in a predominantly verbal society, it might be possible to devise suitable training procedures for the strengthening and control of visual memory as a voluntary option in educated adults. But if, on the other hand, visually gifted individuals should prove to have unique biological endowment, it seems reasonable to ask whether computers can be made to function as prosthetic devices which would compensate for inadequate or missing perceptual or mnemonic mechanisms, on the one hand, or disable or suppress normally-active imagery-inhibiting mechanisms, on the other hand, so that visual imagery can be made to be unusually vivid, persistent, and retrievable. Of course, since the uncontrolled perseveration of visual images is more likely to be distracting rather than useful, the matter of volitional control is of considerable importance. Another possibility that deserves consideration is that, with computerized assistance, even relative strong visual imagery could be further enhanced, thereby extending the limits of human performance. We can say with some certainty that one of the principal barriers to progress in the design of closely-coupled

man-machine systems is our limited knowledge of human memory mechanisms. It seems reasonable to assume that properly designed closely-coupled man-machine systems can be used not only for obtaining superior measurements of human perceptual and mnemonic capacities, but also the enhancement of perceptual and mnemonic skills.

Our plan was to design and implement computerized techniques for real-time monitoring and prediction of visual imagery-relevant brain activities through the use of eye-movements and EEG signals as estimators of these central activities. This predictive information was to be used to control visual displays in a manner designed to enhance visual mnemonic effects. In this report we shall describe our successes and our failures, our progress and our detours, and our revised views of the problems involved in achieving a biocybernetic enhancement of visual perception and memory.

PERSEUS: A STATE-OF-THE-ART EYE-COUPLED SCENE GENERATOR

Accurate real-time monitoring of human eye-pointing constitutes a very considerable technical task. Our advanced computerized eye-tracking system, known as PERSEUS, incorporates the 2-dimensional double-Purkinje-image eye-tracker (DPIET) described by Cornsweet and Crane (1973) and updated by Crane and Steele (unpublished report, Stanford Research Institute). The DPIET is a hardware system consisting of an infrared light source, servo-controlled optics, and associated electronic circuitry. This device tracks and compares the positions of the first and fourth Purkinje images formed by reflections of infrared light beamed at the subject's eye. Inasmuch as these two Purkinje images change their separation when the eye rotates but move the same amount in the same direction when the eye translates, their amount of separation provides a sensitive measure of eye-rotation which is uncontaminated by the principal source of eye-tracking error, viz., eye-translation. However, the DPIET has certain limitations which must be compensated for by the computerized system. For instance, the DPIET knows nothing about higher order phenomena such as fixations and saccades; these must be discriminated as temporal patterns appearing in the continuous output voltages representing

horizontal and vertical eye-positions. Nor does the DPIET have any provision for creating and controlling the display of visual stimuli. Finally, there are the idiosyncrasies of each subject's eye which can produce nonlinear distortions in the eye-tracker estimates of eye-pointing coordinates. In PERSEUS the analysis of fixations and saccades, the generation and control of scenic stimuli, and the linearization of each subject's nonlinearities are all handled by a medium-sized disk-oriented digital computer, the PDP-15/76. This highly accurate (less than 4 minutes of arc error over a field 20 degrees in diameter) eye-tracking system has been used to implement various biocybernetic schemes for controlling scenic displays on the basis of eye-position. It is capable, for example, of stabilizing a scenic stimulus upon the subject's retina without any attachments to the eye. While in this mode it can blank any portion of the display, e.g., periphery, parafovea, or fovea. PERSEUS performs real-time analyses on sequential fixations and saccades, numbers scanpath fixations, and superimposes this information on an ancillary CRT display of the visual stimulus. The data collected during experiments are stored and subjected to more complex forms of off-line analysis. Of particular interest are the visual contact probability plots; these 3-dimensional surfaces are formed by weighting the area surrounding each fixation by an estimate of the local retinal acuity. For quantitative purposes the contact probability surfaces are sliced at equal amplitude intervals to produce a set of iso-contours which are then superimposed upon the scenic stimulus

pattern. By comparing these iso-contours with a similar analysis of the spatial frequencies in the stimulus it is possible to obtain a measure of selective attention which takes into account the eye's natural preference for high frequency "hot spots."

LIMITATIONS INHERENT IN HARDWARE INSTRUMENTS

Sooner or later, the user of a ready-made instrument must come to terms with the capabilities and the limitations inherent in his particular instrument. Efficient exploitation of an instrument requires a good understanding of the instrument's input requirements and its output limitations. It is the user's responsibility to do the work necessary for the proper application of the instrument. The maker of the instrument is only responsible for the intrinsic design and construction of the instrument *per se*. The instrument-maker only agrees to do a limited amount of the user's work in solving certain kinds of problems. Despite all of this, there is a tendency on the part of instrument-users to expect more help from the instrument-maker than the instrument-maker does, in fact, provide. The instrument-user must take up where the instrument-maker leaves off. He must learn to use the instrument properly and, when necessary, he must undertake the task of compensating for the hardware limitations.

The double-Purkinje-image eye-tracker used in this project is no exception to the above-mentioned rule. The

The eye-tracker is designed to operate within certain limits. The DPIET, for example, has a limited angular range over which it can stay in tracking mode. Basically, this limitation is set by the instrument's need to track the first Purkinje image and to track the fourth Purkinje image in relation to the first. Anything that interferes with these requirements will cause the instrument either to lose track and so report or to lose track but continue to report itself in tracking mode (i.e., to track a spurious image). It is the user's responsibility to arrange conditions which will increase the probability that the instrument will stay in the tracking mode and which will decrease the probability of spurious "tracking" responses.

At the output end of the eye-tracker, the instrument makes no provision for recognizing when the eye is "fixating" or "saccading." Therefore, if this information is required by the user, he will have to do the work necessary for accomplishing this further classification of the vertical and horizontal eye-position data. And, if his need for this information is constrained by time requirements, he may have to devise means of accomplishing this classification at high speed. Needless to say, as the performance requirements are increased, the difficulties proliferate.

In our biocybernetic project, we set as our goal the achievement of eye-coupled control of scenic displays which would permit a scenic tartet to be "stabilized" on a subject's retina in a specified location while allowing the subject a

certain amount of freedom in eye-movements. Furthermore, the scenic content must be controllable on the basis of contingent eye-position. That is, without restraining eye-movements (within certain limits) we wanted to be able to modify scenic content in a specified manner on the basis of current eye-position and, if possible, on the basis of anticipated eye-position. Whereas a contact lens-mounted mirror can achieve image stabilization, the information concerning eye-position is not made available for controlling the scene-generator on the basis of eye-position; that is, since the mirror moves with the eye, the image reflected off the mirror will also move with the eye and therefore image-position on the retina will not be influenced by eye-movements. Thus, by means of this physical attachment of the image source (mirror) to the eye and the anatomical attachment of the retina to the eye globe, the image-retina from moment to moment is virtually constant. This method does not compensate for image-retina displacements; instead, it eliminates their source by preventing the slippage from occurring. This approach eliminates the need for tracking and compensating (and for processing information essential for successful tracking and compensation) but it requires an attachment to the subject's eye, on the one hand, and leaves the experimenter with the large and unsolved task of assessing eye-position. And without eye-position information it is not possible to control position-conditional stimulation. Since we were primarily concerned with cybernetic modeling of brain functions, and since the visual system does

not solve its target-tracking problems by means of mechanical coupling of the target to the eye globe, we decided to approach the problem of visual tracking in a manner roughly analogous to that of the human visual system. That is, since we know that the normal human subject tracks a visual target by detecting and measuring image displacement and by making compensatory movements to move the target to a preferred retinal location, we decided to play a similar game, namely, to monitor changes in eye-position and to move the stimulus in a manner calculated to control the placement of the optical image on the retina regardless of the subject's eye-movements. If successful in this endeavor, we expected to reap certain rewards in the form of having a unique capacity to either assist the brain in its tracking efforts or to hamper it. On the one hand, by assisting the brain in its tracking efforts we might be able to extend the performance limits characteristic of the unaided visual system; on the other hand, by interfering in well-defined ways with the subject's tracking responses we might be able to discover some of the brain's tracking strategies which are not transparent to ordinary methods of experimental analysis.

A secondary aim, but one that is essential to developing a sustained assault on the limits of biocybernetic control, was to devise a computerized eye-movement analyzing and display-controlling system that would allow control capability to be continually expanded. By contrast, systems which are inherently capable of doing only one thing at a

time and which preclude the possibility of simultaneous processing and control of more than one variable; this means that in order to mount an attack on one front you must retreat on another front. For example, if you had only the 2-dimensional eye-tracker and the Cornsweet-Crane optometer, you have to decide whether you wish to measure eye-position or visual accommodation because you could only use one instrument at a time (although the optometer has some capacity for measuring eye-position, it is rudimentary as compared with the eye-tracking capacity of the eye-tracker). The recently developed 3-dimensional eye-tracker (Crane and Steele, unpublished report) overcomes this limitation and permits the user to track horizontal and vertical eye-position while simultaneously tracking accommodation. As we shall see, this increase in the capacity of the instrument, while it removes certain experimental limitations, poses many problems for the user who desires to exploit the possibilities of the instrument.

THE DOUBLE-PURKINJE-IMAGE EYE-TRACKER (DPIET)

The two-dimensional eye-tracker developed at the Stanford Research Institute (Cornsweet and Crane, 1973; Crane and Steele, 1976) utilizes an advanced method of eye-tracking which is based on servo-controlled tracking of two Purkinje images, namely, the first Purkinje image and the fourth Purkinje image, which move identically during translational movements of the eye but which move differentially during rotational movements of the eye.

The virtual image resulting from light reflected at the interface between the air and from the cornea is known as the first Purkinje image. The second Purkinje image, formed by light reflected at the interface between the rear surface of the cornea and the aqueous humor of the anterior chamber of the eye, is practically coincident with the first Purkinje image. The third Purkinje image, a virtual image, is formed by light reflected at the interface between the front surface of the lens and the aqueous humor; this image is much larger and more diffuse than the other images and is formed, fortunately, in a plane which is remote from the plane of the other Purkinje images. The fourth Purkinje image is formed by light reflected at the interface between the rear surface of the lens and the vitreous humor that fills the posterior chamber of the eye globe. The rear surface of the lens functions as a concave mirror; its reflected light forms a real image of the light source.

The first and fourth Purkinje images are almost the same size, and are formed in almost exactly the same plane. Due to the fact that the change in the index of refraction at the lens-vitreous interface is much smaller than that at the air-cornea interface, the intensity of the fourth Purkinje image is relatively weak, i.e., less than 1% of the intensity of the first Purkinje image.

During translatory movement of the eye the first and fourth Purkinje images move in the same direction and travel the same distance (i.e., the distance between them does not

change). However, when the eye rotates, the first and fourth Purkinje images move in opposite directions, causing them to move either closer together or farther apart. The distance separating the first and fourth Purkinje images thus provides a basis for measuring the angular rotation of the eye without the error contributed by translatory movements of the eye.

The 2-dimensional double-Purkinje-image eye-tracker provides two separate output voltages (continuous) whose magnitudes are proportional to the horizontal and vertical components of the distance separating the first and fourth Purkinje images formed in the eye in response to the input of an infrared light beam. This 2-dimensional DPIET provides continuous monitoring of the angular position of the subject's eye with a high level of accuracy and without the need for mechanical contact with the eye. Since the eye-tracker operates in the infrared, it does not interfere with normal vision.

The double-Purkinje-image eye-tracker (DPIET) is limited in range by several factors. First, when the eye rotates toward the infrared input axis (i.e., to the right), the horizontal distance between the first and fourth Purkinje images decreases. When light from either the first or third Purkinje images invades the photocell which is supposed to see only the fourth Purkinje image, it becomes impossible to continue to track the fourth Purkinje image. Second, when the eye rotates away from the infrared input axis, the

horizontal distance between the first and fourth Purkinje images increases. Tracking in this direction becomes impossible when the iris cuts off the fourth Purkinje image. Third, vertical rotation of the eye causes corresponding vertical separation of the first and fourth Purkinje images; tracking of vertical rotation ceases when the iris cuts off the fourth Purkinje image. In brief, the range of the DPIET is limited by the pupil boundary in rotational directions away from the infrared input axis and by lights from the first and third Purkinje images confusing the fourth Purkinje image detector when they move too close to the fourth Purkinje image as the eye rotates in the direction of the infrared input axis. From this it will be evident that pupil size is an important limiter of tracking range; the larger the pupil size, the greater the tracking range in all directions except toward the input axis. Therefore, since reduced ambient light favors a larger pupillary size, it also increases the tracking range of the DPIET.

To minimize tracking loss during eye blinks, the gain of each servo motor driver is sharply reduced during each blink. Blinks are detected by summing the outputs from the four quadrants of the first Purkinje photodetector. This sum, which is nominally independent of eye direction, is constant except during eye blinks; during a blink this sum will increase or decrease substantially depending upon the amount of light reflected from the eyelid. The blink detection circuit indicates that a blink is in progress (output

voltage change) when the sum is outside the adjustable upper and lower limit values. To minimize tracking loss during eye blinks (that is, to prevent the tracking servo motors from moving arbitrarily) the gains of the servo motors are sharply reduced during each blink. If the post-blink eye-position does not differ greatly from the pre-blink eye-position, the servo controls can quickly resume tracking their respective Purkinje images when the eyelid opens.

Our updated model of the DPIET (Crane and Steele, unpublished report) is designed to tolerate up to one centimeter of variation in eye position in horizontal, vertical, and axial dimensions. The axial tolerance is achieved by incorporating automatic focusing. The horizontal-vertical tolerance is achieved by making the input light path track eye position automatically.

The first and fourth Purkinje image trackers are both provided with output signals indicating when they are unlocked (i.e., out of tracking mode). These signals provide a convenient means of assessing eye-tracking records. We have incorporated them into our system for mapping the trackable perimeter.

PERIMETER

In order to exploit the full range of the eye-tracker it is necessary to measure the limits of eye-tracking in every direction. We have developed a routine for automatic mapping of the trackable perimeter. This routine, called

PERIMETER, allows the experimenter to select up to 24 radial lines arranged around a central field point. The subject is instructed to track the movement of the CRT beam as it moves centrifugally at a constant speed (selectable by experimenter) along one of the radial lines. The peripheral tracking limit is marked when the fourth Purkinje tracker unlocks. Then the subject returns to the central fixation point and the target beam moves out along another of the radial lines. And so forth. See Fig. 1a. When all of the selected radii have been tracked to the limit of trackability, the program then connects the adjacent radial limits to form a trackable perimeter. Each of the radial limits is measured to the nearest tenth of a degree and this value is displayed on the graphic perimeter. See Fig. 1b.

This method of mapping the trackable perimeter allows the experimenter to adjust his stimulus materials so that they fall within the trackable limits. This means (in the digital domain) that a picture can be offset vertically and horizontally (with respect to the central fixation point) until it is centered in the trackable space and that the scale of the picture can be increased so as to fill the trackable area or reduced so as to fit inside the trackable area. Without accurate measurements of the trackable perimeter, the experimenter is inclined to make two types of error: first, in order to keep his picture well inside the tracking range of the eye-tracker, he may make his pictures smaller than necessary, an excessively costly form of

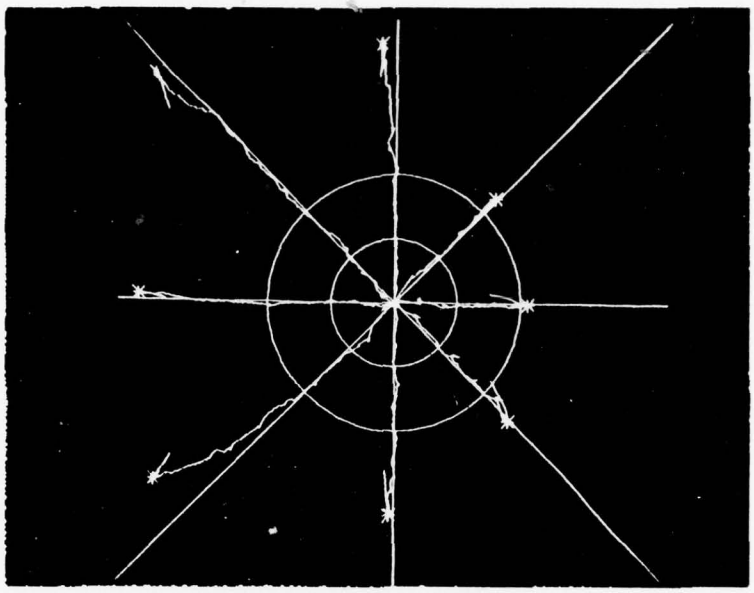

FIGURE 1a. Perimetry data obtained from subject J.N. The eye-tracker monitors the subject's eye as it pursues a target that moves outward along each of the radii. The star marks the point at which the eye-tracker loses track, i.e., the trackable limit.

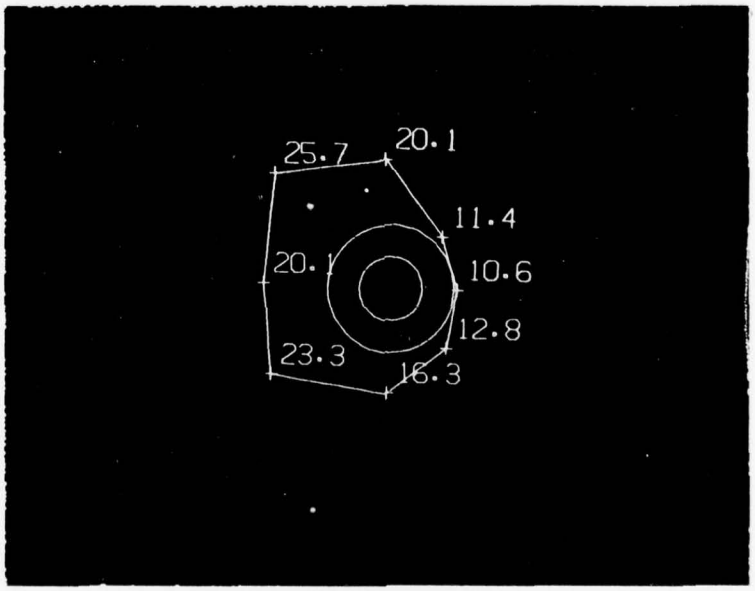

FIGURE 1b. The perimeter of trackability automatically computed from the data in Fig. 1a. The circles are spaced 5 degrees apart.

insurance against losing track of eye movements. Second, he may make his picture so large that parts of it lie outside the perimeter and eye movements directed at those parts are untrackable. On the one hand, it is desirable to make the picture as large as possible in order to obtain a maximum angular separation between features; on the other hand, such expansion runs the risk of making parts of the picture fall outside the range of trackability. It is a case of the perennial trade-off between field size and resolution. Both the human eye and the eye-tracking machine have limited fields and limited resolving powers.

LINEARIZATION OF EYE-TRACKER ESTIMATES

We notices that the response of the eye-tracker to the eye-movements of a real eye was subject to nonlinear distortions, i.e., the distortions were different in different parts of the trackable space. Mind you, these nonlinearities are not large when compared with other eye-tracking devices but they significantly limit the functional resolving power of the eye-tracker for an extended range of movements. The actual human eye is, after all, not the idealized form anticipated in the design of the hardware. The distortions observed are nonlinear in the sense that no single set of vertical and horizontal gain and offset adjustments will correct the errors in all parts of trackable space. What is required is local control of gains and offsets, i.e., independent gain and offset controls for each part of the trackable space. By linearization of these nonlinear distortions

we could expect to extend the spatial resolving power of our eye-tracking system.

For the above-mentioned purpose we have implemented an adaptive filtering scheme to obtain the required corrections. In our procedure we ask the subject to fixate sequentially each calibration dot in a matrix of 11 x 11 dots; each calibration dot is located two degrees away from its nearest vertical and horizontal neighbors. Actually, the computer first displays the entire matrix of calibration dots (in the memory mode of the CRT) and then brightens the particular dot that the subject is currently required to fixate. When the subject is satisfied with his fixation of that dot, he presses a button which causes the computer to sample the X and Y outputs of the eye-tracker. Then the computer immediately brightens the next dot to be fixated, and so forth. It is assumed that the subject is able to properly fixate the calibration dot and that an adequate selected sample of the tracker output will properly estimate this fixation, but subject to the nonlinear distortions we are now considering. When the subject has completed his fixation of all of the calibration dots, the computer immediately displays the matrix of calibration dots and superimposes the fixation points as estimated by the eye-tracker. The task of the corrective filter is to adjust the eye-tracker estimates so as to minimize the discrepancies between the known fixation points (i.e., the calibration dots) and the corresponding eye-tracker estimates.

The developing filters are graphically displayed in Figs. 2a and 2b. Figure 1a illustrates the status of the filters prior to any adaption. The X-filter surface is located in the lower left portion of these figures; the Y-filter, in the lower right portion. The filter is represented by a matrix of 21 x 21 dots (twice the density of the calibration dots) and is depicted as lying in a horizontal plane. This plane represents zero correction; displacements above this plane represent corrective additions; displacements below this plane, corrective subtractions. The upper left portions of Figs. 2a and 2b show the calibration points (11 x 11), the eye-tracker estimates of fixations, and ellipses indicative of the magnitude of the error. In the upper right portions of these figures are displayed three useful statistics of error; these are updated after each adaptive iteration (an adaptive scan of all the data samples). The uppermost statistic is the maximum error expressed as a percentage of the viewing angle (in this case it is 20 degrees). The second statistic is the mean square error expressed as a percentage of the viewing angle. The third statistic is similar to the second but expressed in terms of CRT (scope) display units (there being 1024 units across the display). Figure 2b illustrates the condition of these displays and statistics after the adaptive filtering process has greatly reduced the nonlinear errors and the filter surfaces have approximated their asymptotic values.

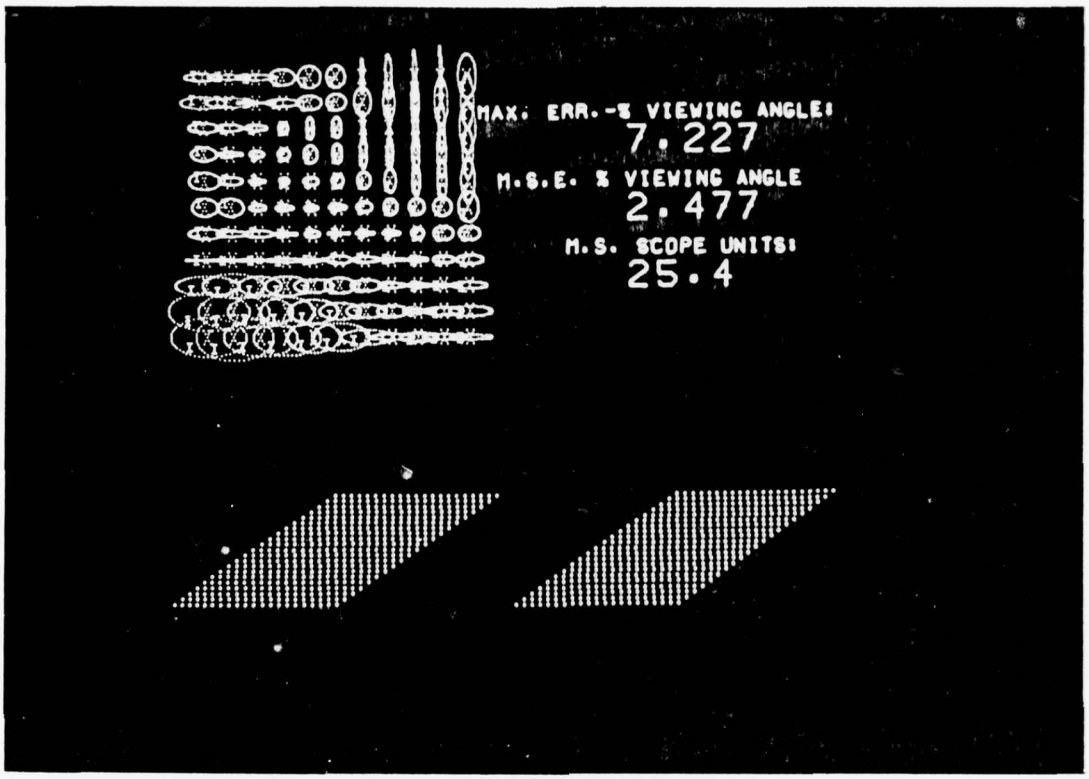

FIGURE 2a.

Calibration data and adaptive filter surfaces prior to adaptive correction. The 11 X 11 calibration matrix is shown in the upper left portion of the figure. The ellipses are indicative of the magnitude and direction of the errors to be corrected. The zero-correction filter surfaces for horizontal and vertical components are shown in the lower left and lower right portions respectively; they are depicted in pseudoperspective as horizontal planes. For purposes of interpolation there are twice as many filter points as calibration points. Various useful statistics are shown in the upper right portion of the figure. First, there is "maximum error" expressed as a percentage of the viewing angle (20 degrees in this case); second, there is "mean square error" expressed as a percentage of the viewing angle; third, there is "mean square error" expressed in scope units (1024 units X 1024 units). After each filter scan, the filter surfaces are adjusted, the error ellipses are reduced a corresponding amount, and the error statistics are updated.

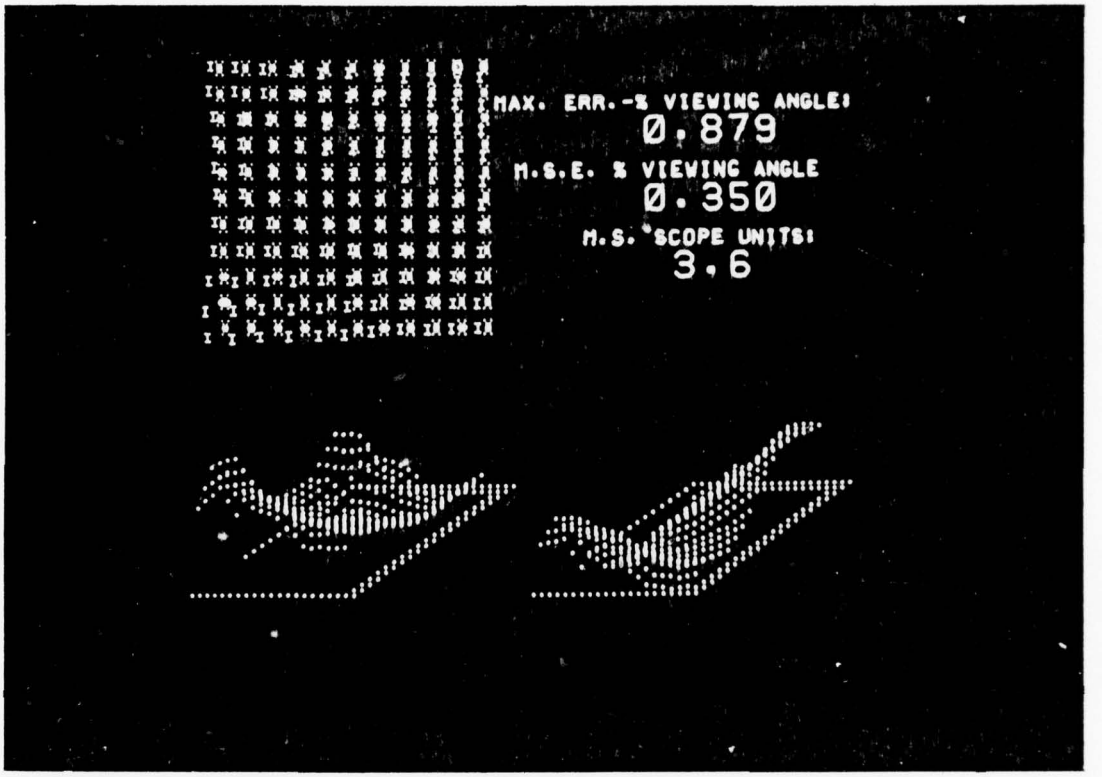

FIGURE 2b.

Calibration data and adaptive filter surfaces after adaptive correction is approximately asymptotic. Displacement of the filter surface upward from the zero-correction plane (defined by edge values) indicates the amount which must be added to the eye-tracker output to correct for nonlinear distortion; conversely, displacement of the filter surface downward from the zero-correction plane indicates the amount which must be subtracted from the eye-tracker output to correct for nonlinear distortion. In PERSEUS these corrections are automatically made in the eye-tracker data by interpolation from the nearest block of 5 filter points. The inverse of the corrective filter can be used to distort a digital picture in order to obtain an intuitive impression of the 'cost' of neglecting to correct the nonlinear errors.

These filters are somewhat analogous to optical lenses designed to correct for astigmatic visual distortions. We have observed that the filter characteristics obtained on different days from the same subject are surprisingly similar. Stated otherwise, this means that the filters are quasi-stationary or relatively invariant with changes in time. Thus, rather than starting the calibration process "from scratch" on each separate experimental occasion, it is possible to install a previously-generated filter, adjust the offsets of its central point, and adapt from there. In the terminology of control theory, this usage of relevant prior information is called a feedforward operation because it allows the system to advance to the "ball park" of the optimal filter. Such exploitation of stationary parameters of a measurement system provides for greater convenience and greater speed in the operation of the system. In some cases exploitation of these properties is essential to successful tracking.

SEER: ROUTINES FOR THE CREATION, EDITING, STORAGE, AND DISPLAY OF DIGITAL PICTURES

In order to gain the full advantages that derive from having a real-time monitor of eye-direction it is desirable to have access to computer-generated visual targets. Then if the computer knows where the eye is pointed, it can move the target into some specified relation to the eye. Because of the great speed of the computer, the visual effect is

equivalent to physically coupling the optical image to the retina. Yet, unlike stabilized images produced by optical systems physically attached to the eye, the digital image is subject to highly controlled manipulation. However, the generation and management of non-trivial digital pictures can be a rather complex matter.

Conversion from analog scenes to digital strings is facilitated by the use of a GRAFPEN system which employs an ultrasound-emitting pen tip and a drawing surface bordered at the top and the left side by ultrasound sensor strips. The GRAFPEN system outputs X and Y voltages which are proportional to horizontal and vertical positions of the grafpen. These voltages pass to the A/D converters of the PDP-15 computer. Although it might seem that this arrangement would solve the problem of converting analog scenes into digital number strings, it proves to be awkward and inefficient. As our needs became apparent, we developed an increasingly elaborate set of routines for making and handling digital pictures. We have used the acronym SEER (Stanford Eye Experiment Routines) to designate the software developed for this purpose.

The work of making and using digital pictures can be divided conveniently into three parts. First, there is the work of the copyist who has to trace the analog picture contours with the grafpen and control the computer's sampling, labelling, and storing operations. Second, there is the task of assembling contour segments into a "picture," editing

it, naming it, and storing it as an entity. Third, there is the task of assembling a sequence of pictures for an experiment. The various SEER routines are designed to facilitate this work. The transfer of these digital pictures from one type of storage to another must also be provided; that is, the digital points have to be transferred from core memory to fixed-head disk, or from fixed-head disk to DECtape or 9-track digital tape, or from fixed-head disk to movable-head disk, or in the reverse direction. The various SEER routines are designed to make these transfers relatively easy.

In most of these activities it is desirable that the experimenter be able to see what he is working with. Therefore, we have provided extensive display capabilities.

Now, to consider the work of the copyist: the scenic material to be copied is rear-projected onto the ground glass surface of the grafpen drawing area. The copyist must then use his eyes to guide his hand as it traces a selected contour with the tip of the grafpen. Since this form of copying is tedious at best, it is convenient to arrange the system so that he has flexible command of the computer. The copyist's dialogue with the computer is carried out through the use of the grafpen. A computer command "menu" is provided in the lower right-hand corner of the grafpen drawing surface. By touching the grafpen to the various designated areas, the copyist can request that the computer give him an element number; the computer's selected number is then displayed

by the computer on a display CRT located near the drawing surface. In the MENU-mode the copyist can draw, erase, label, and store up to 1000 picture segments at a time.

We shall not elaborate on the available menu commands other than to mention that in addition to the more obvious computer commands we have provided some "geometrical assistance" commands which allow the copyist to enlist computer assistance in drawing straight lines, circles, arcs, and rectangles. Also, to economize on digital storage and to obtain an approximately constant density of sample points from a picture contour, there is a spatial sampling routine which stores a new data-point only if it is located some specified distance away from the previously-stored data-point.

Equalizing the sample-density of the picture contours not only economizes on storage but it lays the foundation for meaningful manipulation of contour density. Thus, for example, it is possible to get reasonably good tracings by displaying every other point, or some suitable fixed ratio of points. Or, if you wish to generate a moving target point for the eye to follow, a fixed rate of display of the points will appear to speed up in approaching curves and corners and will appear to slow down in the straight sections. By increasing the sample density in curves and corners it is possible to equalize the apparent speed of the moving light beam on the display CRT. Since we could not arrive at any generally satisfactory formula that would

produce subjective equality of trace-speed, we hit upon an empirical solution whereby the editor is permitted to select any particular segment by marking two limiting points on the contour; then he can draw with the grafpen whatever function seems reasonable and the density-adjusting routine will increase or decrease the sample density along that segment in accordance with the selected function. Then the editor can request that the segment be traced on the CRT using the adjusted density. If he is dissatisfied with the result, he can modify the density function and see if that produces a better subjective effect; if he is satisfied with the result, he can move on to the next segment. Finally, we should mention that the equal-density contour drawings provide the basis for generating the spatial density matrices for these pictures (to be described later).

Having collected a set of picture segments, the copyist can then proceed with the task of assembling "pictures" from picture-segments. He can edit his picture in many ways. For instance, he can magnify it or minify it, he can offset it on horizontal and vertical axes, and he can rotate it (yaw, pitch, roll). Then he can give this edited version of the picture a label before storing it in one of several convenient places.

Finally, he can assemble a set of "pictures" into an experimental sequence, specifying various display parameters for the experiment. With these manipulative capabilities it is possible to utilize digital pictures much the way that

one would use a set of photographic slides and a projector. However, being in the digital domain, the digital picture can be manipulated by the computer in some extraordinary ways. On the negative side, there is the restriction (in our set-up, at least) to monochromatic contour-pictures or "cartoons."

As a future development, we foresee the coordination of visually-rich color T.V.-type of display with the schematic contour-pictures, allowing photographs to be overlaid by corresponding schematics on a frame-by-frame basis.

FIXATION-CONDITIONAL STIMULUS PRESENTATION

In many experiments, the subject is asked to fixate a centrally located target as a pre-condition for the exposure of the stimulus pattern. From the point of view of experimental design, fixation of the central point is a stimulus parameter in the sense that the specified fixation arranges for consistent retinal placement of the stimulus pattern. Yet it is seldom that any precautions are taken to monitor or control this parameter other than to instruct the subject to fixate and expect his cooperation. A more rigorous approach would be to arrange for electronic monitoring of eye-pointing and for electronic circuitry for making stimulus presentation conditional upon fixating within prescribed spatial boundaries for some prescribed period of time. This arrangement has the advantage that the stimulus exposure can be extended as long as the subject keeps his eye within the

prescribed spatial boundaries; the alternative, of course, is to employ tachistoscopic exposures which are sufficiently brief as to preclude significant reactive eye-movements. A computerized fixation-conditional stimulus controller is illustrated in Figs. 3 and 4. This is actually the simplest use of a more general system of visual scanpath analysis which is described below.

REAL-TIME SCANPATH ANALYSIS

Automatic methods for the off-line analysis of scanpaths are described in Shah (1977). We shall now describe a real-time method for the analysis of scanpaths. It is evident that an effective real-time scanpath analyzer is an essential development in the process of implementing cybernetic models for the prediction and control of ongoing eye-pointing behavior.

This computer program was first implemented at Ames Research Center on an SEL-840 computer which has an Evans and Sutherland display system. Two large CRTs are utilized: (1) a monitor scope on which are displayed the various analyses that are of interest to the experimenter, and (2) a stimulus display CRT which is viewed by the experimental subject. During the past year we have developed an equivalent real-time scanpath program for our PDP-15/76 at Stanford.

The subject is seated in front of the stimulus CRT, his head stabilized by a bitebar, with the distance between the eye and the CRT surface adjusted so that 12 inches in

either the vertical or horizontal dimension on the scope corresponds to 20° of visual angle. His right eye is monitored by a 2-dimensional eye-tracker which provides output voltages proportional to horizontal and vertical eye-direction.

Following completion of suitable calibration procedures, the subject is presented with a cartoon-type stimulus (out of digital storage) on his display scope and instructed to examine it freely for a specified number of seconds or for a specified number of fixations, after which the cartoon disappears and the central fixation spot is restored. While the subject is examining the cartoon, the X and Y voltages from the eye monitor are entered into X and Y voltage distributions located below and to the left, respectively, of the fixation circle (see Fig. 3). The entries in these distributions are cumulated in accordance with a sliding time-window (the width of the window being specified in milliseconds by the experimenter). The peak of each distribution (X and Y) is treated as the best estimate of the eye-pointing during the time constant characterizing the width of the sliding window. This X,Y location is made by the center of a circle (solid line) the radius of which is controllable by the experimenter. We sometimes refer to this circle as a "fixation circle" or as a "target circle" depending upon the experimental use being made of it. If we set minimum amplitude criteria for the distribution peaks, we have a basis for monitoring the "swell time" or "spatial

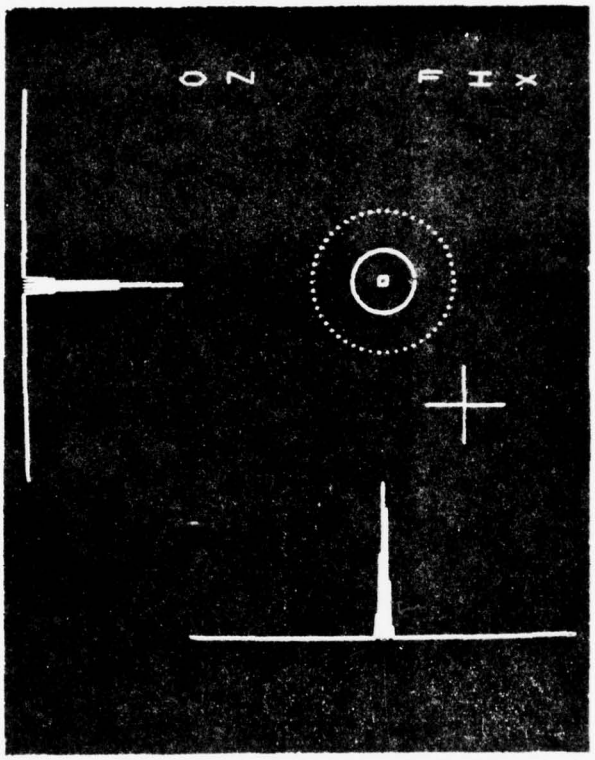

FIGURE 3

Real-time eye-fixation monitor

CRT display of horizontal and vertical eye-direction distributions. Instantaneous eye-direction is indicated by the small square located inside the solid circle (or target circle). The variance limit is designated by the dotted circle; a prior fixation, by the cross. "ON" indicates that the point-of-regard (small square) has remained inside the solid circle long enough to meet the specified requirement for being "on target". "FIX" indicates that the system is in the "fixation-display" mode and will collect and mark the locations of every different fixation-location.

invariance" as a function of time, this is a reasonably effective way of defining the beginning of each fixation. The method for defining the end of a fixation is described in the next paragraph.

A second circle, sharing the same center as the target circle, has an independently variable radius which is equal to or greater than that of the target circle. This circle, which appears on the monitor scope as a dotted line (see Fig. 3), is used to set a spatial variance limit around the target circle. When the instantaneous eye-position exceeds the perimeter of the variance circle, it is assumed that the eye is making a saccadic movement, or at least, that it is moving sufficiently far away that the information accumulated in the distributions is irrelevant, so the distributions are re-set to zero and the existing target circle vanishes. When a new fixation is detected, a new target circle and a new variance circle will appear. And so forth.

Another important feature of this program is the arrangement whereby if the instantaneous eye position (the small square inside the target circle in Fig. 3) remains inside the target circle for a specified number of milliseconds (determined by the experimenter), a voltage appears on an appropriate analog output line and the word ON is displayed on the monitor scope until the instantaneous eye position travels outside the target circle; when the output voltage goes to ground the word OFF replaces the word ON (see Fig. 4).

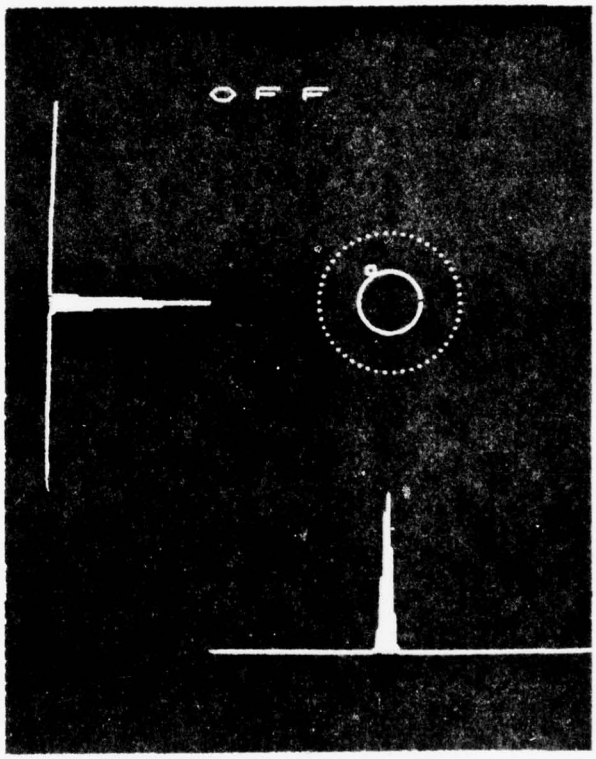

FIGURE 4

Real-time eye-fixation monitor

Similar to Fig. 3 except that the point-of-regard (small square) is now just outside the target circle. "OFF" indicates that the subject's eye is off-target; this means that the visual stimulus would be withheld in the fixation-conditional stimulus-presentation experiment described in the previous section. If the point-of-regard were to stray outside the established variance limit (dotted circle), the current fixation would be considered to be terminated and the system would abort the eye-direction distributions and start collecting a wholly new set on the assumption that the eye is saccading to a new fixation-location.

This ON/OFF feature provides the basis for the fixation-conditional visual stimulation described earlier. The real-time scanpath program provides for the collection of up to 100 fixation-locations before it must pause momentarily to transfer these data to another storage area. In the scanpath display, a cross appears at each fixation-location and each cross is numbered in accordance with its ordinal position in the scanpath.

Figure 5a and Fig. 5b contain a cartoon of a painting which depicts Judith with the head of Holofernes. In Fig. 5a the current location of the point-of-regard is marked by the small circle at the tip of Holofernes' nose. This eye-position is indicated by the X-probability and Y-probability distributions located below and to the left, respectively. When these probability distributions reach some selected peak amplitude, the computer recognizes that location as a fixation and deposits a small cross in the display to mark that spot. Of course, information concerning the location, duration, and scanpath number is stored for subsequent use.

We have employed a number of different techniques for the analysis of scanpath data. One of the current developments is illustrated in Figs. 6a and 6b. Although the usual endpoint of fixation-detections is reached when fixation locations are superimposed on the scenic target, it is possible to consider more complex transforms. In this approach, we have tried to give recognition to the fact that some information is processed through every part of the retina, not just

FIGURE 5a

Judith with the head of Holofernes

The eye-direction distributions (below,and to the left of the picture) indicate that the subject's point-of-regard is near the tip of Holofernes' nose. Prior fixations are marked on Holofernes' eyes. This is the type of real-time display that PERSEUS presents on the experimenter's CRT; normally, only the picture is displayed on the subject's CRT.

FIGURE 5b

Judith with the head of Holofernes

The picture is now covered with a sizeable number of fixations which are marked by small crosses. Notice that the subject has a predilection for faces (eyes and mouths, in particular) and gives considerable attention to Judith's bosom (no doubt as the artist intended for he has placed her bosom in the center of his compositional circle).

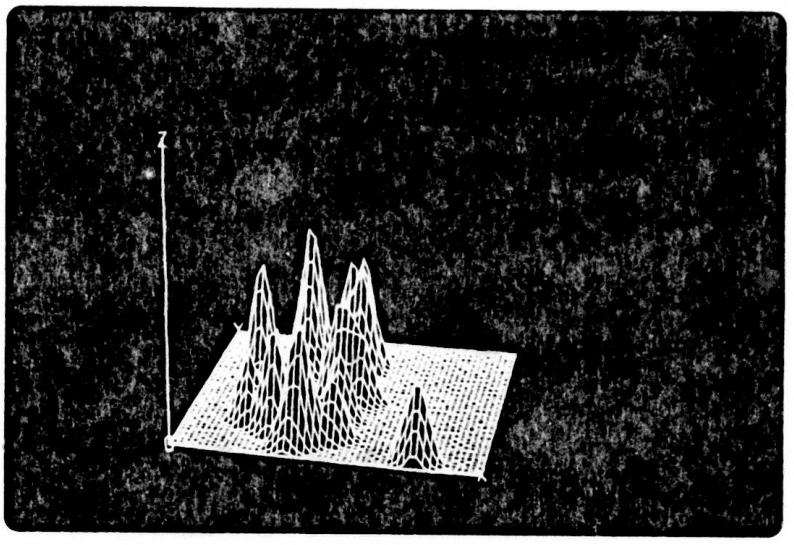

FIGURE 6a. Three-dimensional visual contact probability matrix based on the fixation data shown in Fig. 5b. The picture plane is horizontal.

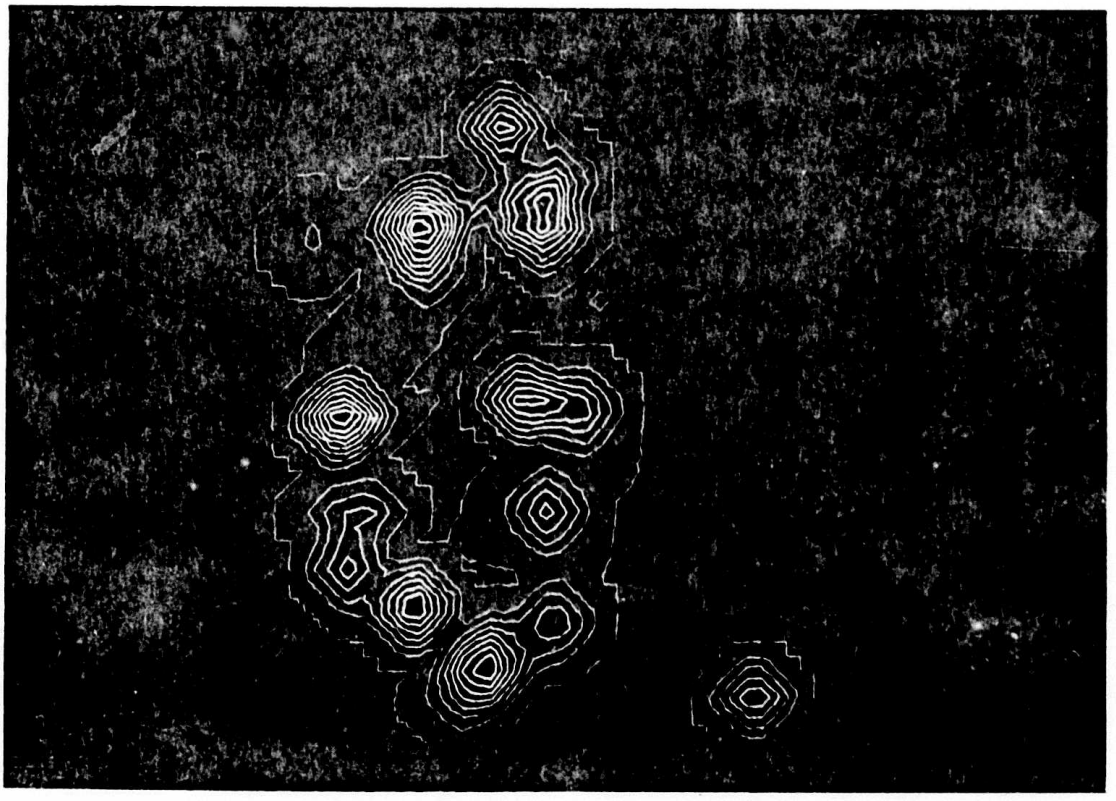

FIGURE 6b. Iso-contours cut from the matrix in Fig. 6a by equally-spaced planes parallel to the picture plane. The scale is increased to match the picture in Fig. 7a.

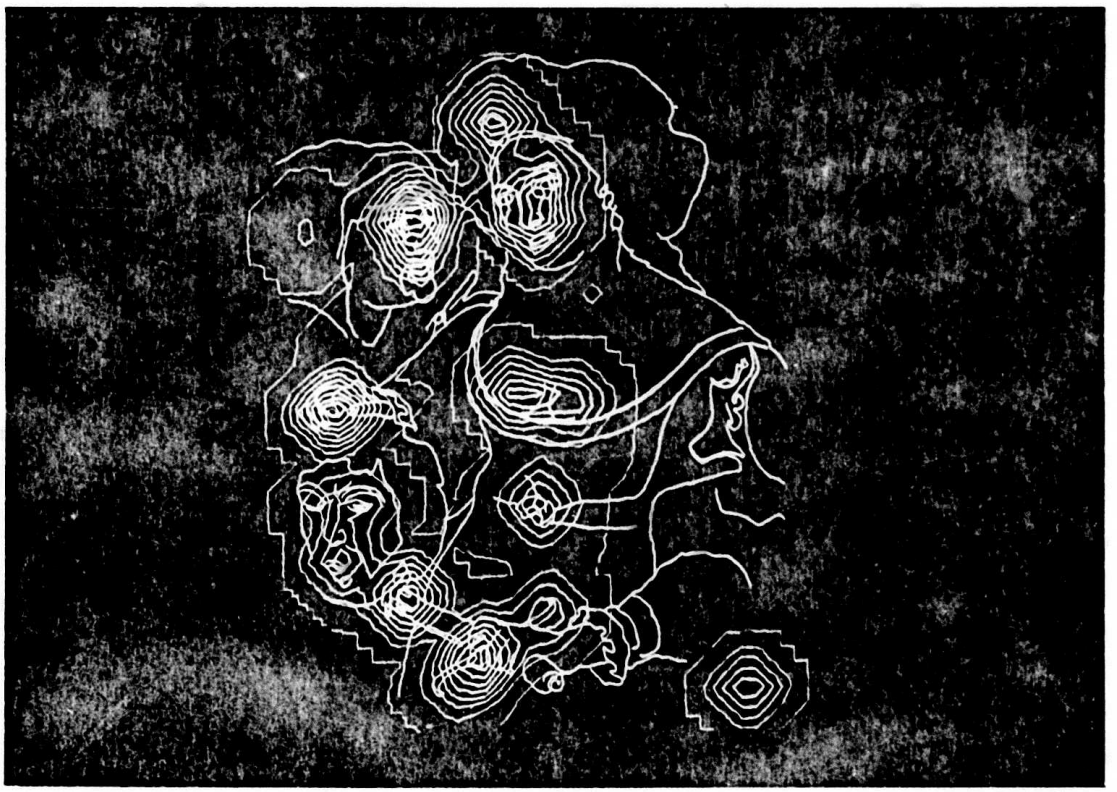

FIGURE 7

The iso-contours of Fig. 6b superimposed on the picture of Judith with the head of Holofernes. Obviously, the spread of visual contact around the centers of attention will depend upon the shape of the contact probability weights entered around each fixation point. This system is flexible in that it allows the experimenter to select his weighting pattern on whatever basis he deems advisable. Presumably, the pattern of contact probability changes depending upon environmental illumination, subject motivation, etc. We would suppose that the retinal acuity function provides good basic weighting which could be modified on the basis of other considerations.

through the foveal portion. In this case we have made a weighted entry into every portion of the picture space surrounding the fixation point, based upon some acuity estimate. The exact form of the acuity weighting is immaterial for the current exercise; the program permits the entry of any specified set of weights. What this form of display helps to formulate is some concept which we might call "visual contact probability." Figure 6b illustrates a sample probability matrix (simulated 3-dimensional plot). Figure 6a shows (enlarged) the equal amplitude transform (or "geophysical map") of the data in Fig. 6b. Figure 7 illustrates such an equal amplitude transform of the contact probability matrix for a set of fixation obtained from a subject who was given an opportunity to examine this scene. The equal amplitude contours are projected onto the cartoon (Judith with the head of Holofernes). /Note: the lower right set of contours--off the picture--resulted from the fact that the subject left the bitebar before sampling was terminated; that "fixation" is the default location of the tracking system in the absence of a trackable eye./ A dynamic version of this contact probability display provides a view of the contact probability for a relatively narrow sliding window (of time) but this display is more suited to motion picture representation.

MODELING SACCADIC EYE-MOVEMENTS

A number of models have been proposed in the literature for the saccadic control system. While there is a

general agreement regarding the various characteristics discussed above, there is no single model which explains them all.

A natural division of the system is to have a controller driving a plant. The controller is the brain with the target signal as input. The nerve signals produced by this controller drive the plant which is composed of the eye muscle assembly. Models have been proposed for both parts. Westheimer (1959) suggested a second-order linear model for the plant. When such a system is driven by a step input, a saccade-like trajectory results. In Westheimer's model, the plant is underdamped. The model parameters have no physiological significance. Yarbus (1967) suggested that the saccadic velocity curve is sinusoidal. This, however, is an over-simplification. Based on the 200 ms reaction time to a pulse target movement, Young and Stark (1962) proposed a sampled data control system which samples the error and then executes a correcting movement 200 ms later. The target movement in this interval does not cause this initial response to change.

The sampled data model was presented by Young and Stark (1962) and analyzed in detail by Young (1962). It is an elegant application of classical control theory to the study of a physiological system. For the plant, this model used Westheimer's (1959) underdamped second-order system. The model is adequate for simple saccadic responses but cannot predict the correct answer to a pulse-step target. The controller has to be more complex to produce the variety of

responses observed for different inputs. Actually, as Robinson (1973) pointed out, the plant is overdamped and needs a pulse superimposed on a step (a pulse-step) as input to produce the observed saccadic trajectory. Robinson found evidence that the force applied to the extraocular muscles during a saccade is actually a pulse-step. He used data gained through external loading on the eyeball as well as through experiments under normal conditions to determine the parameters of the second-order plant model. Robinson (1973) has suggested a series of models all using the same plant but with controllers of increasing complexity.

Some of the shortcomings of Robinson's models have been pointed out by Cook and Stark (1967). To understand these, we must refer to some neurophysiological conceptions of the saccadic system. Very briefly, the signals from the retina reach area 17 of the occipital lobe via the optic nerve and the lateral geniculate body. The error signal is probably formed in areas 18 and 19 and processed to form the saccadic motors command signal in the frontal eye fields (Fuchs, 1971). This command signal ultimately reaches the oculomotor nuclei. The force applied to the eyeball is a result of the firing of the oculomotor nuclei and the action of the extraocular nuscles. Robinson pointed out that the firing pattern of these nuclei follows a pulse-step, with a high firing rate during the saccade which settles down to a steady value after the saccade. To produce this pulse-step, he proposed that the error signal drives a pulse generator

which produces a pulse, the area under which equals the desired saccade amplitude. This pulse is fed to a network consisting of a neural integrator in parallel with a feed-forward path (the medial longitudinal fasciculus). The output of this network represents the firing rate of the oculomotor neurons. It then drives the second-order over-damped plant (i.e,, the extraocular muscles and attached eyeball) to produce the saccadic eye-movement.

There is considerable physiological evidence indicating the presence of a neural integrator. Robinson (1972) has combined the vestibulo-ocular system, the smooth pursuit system, the saccadic system, and the vergence system into a single model sharing a neural integrator and a feed-forward path. Also as pointed out above, the plant parameters used in his model were derived from physiological experiments. Thus, it is clear that Robinson's model and its parameters have considerable physiological significance, which is essential in a model that purports to explain a biological system.

Robinson's model does have some limitations, however. The eyeball is controlled by 3 pairs of extraocular muscles used for rotation about the X, Y, and Z axes respectively. For example, during a horizontal saccade, one muscle (the agonist) shortens in response to an increase in neural activity whereas the other muscle (the antagonist) lengthens in response to co-responding decrease in neural activity. Cook and Stark (1967) have pointed out that Robinson does

not describe what portion of the driving force is attributed to the agonist and what portion to the antagonist. The model is incomplete in that sense. They also reported that the velocity curves of Robinson's model do not agree with experimental data.

By considering muscle dynamics and the different forces applied to the agonist and the antagonist muscles, Cook and Stark (1967) proposed a nonlinear model for the plant. The model parameters were determined from physiological experiments. They also published comparisons between (a) model-simulated eye-position and velocity curves and (b) experimental data (Cook and Stark, 1968). Clark and Stark (1975) reported that this model produces time-optimal responses.

In its full form, the model proposed by Cook and Stark is a sixth-order nonlinear system and is quite complex; thus it is clear that the saccadic system cannot be fully represented by any simple model. The separation into a controller and a plant is itself a simplification which may not be fully correct. For full understanding of the saccadic system, it is necessary to use more and more intricate models and these will continue to be developed as more becomes known about the system. In general, how complex the model must be depends on its application. Our interest is in monitoring and predicting saccadic eye movements on a real-time basis. Thus, the input-output characteristics of the system are of more interest to us than the exact components. We have to give priority to computational feasibility for real-time monitoring

and prediction. Because Robinson's model is founded on considerable neurophysiological evidence and because of its relative simplicity, we have used it as a starting point for our model development. We have improved upon his model so that we obtain better fit of the model responses with experimental data. Our model parameters are estimated by curve fitting techniques. As a result, we have developed a model that, although it may not be completely sound physiologically, does in fact reproduce the input-output behavior of the saccadic system with sufficient accuracy to be adequate for the purpose of monitoring and prediction.

Let us look at some of the shortcomings of Robinson's model and discuss our modification for remedying the situation. Firstly, it is quite clear that the saccadic system is a nonlinear system, as described before; yet we are using a linear model. This use is justified partly by the fact that the controller, which generates the input to the plant in response to the visual stimulus, is nonlinear, and partly by the need for computational simplicity. The magnitude of the visual stimulus affects both the height and width of the pulse input to the plant so that even though the model itself is linear, velocity saturation and dependence of saccade duration on amplitude are observed. Within the plant itself, the nonlinear muscle dynamics have been approximated by a linear system.

Cook and Stark (1967) pointed out that the velocity curves produced by Robinson's model did not match experimentally-observed velocity curves. In addition, the model as initially

presented does not produce overshoots or undershoots in the saccade. Robinson (1973) suggested a slight change to one of the time constants to produce the desired response. However, we found that such a modification still does not match the observed overshoots. This is a serious defect, particularly since we are interested in the input-output characteristics and not the physiology per se. We have made an important modification to Robinson's model by changing the form of the input to the plant. Interestingly enough, this is the form of the input that would result from Cook and Stark's model if we ignored the differences between the agonist and the antagonist muscles and simply added (algebraically) the forces on the two. This change enables us to accurately fit the model generated position and velocity responses to the experimentally-observed ones for different types of saccades. This improves upon Robinson's model by removing its main distraction.

Until now, most models of the saccadic system have attempted to explain only the gross behavior of the system. Essentially, that means a study of the static behavior of the system, with the primary emphasis being on whether or not a saccade will be produced in response to a specific stimulus. The detailed dynamic characteristics of the saccade have been of secondary importance. Only Cook and Stark have attempted to fit model-generated responses to actual ones. We have essentially used an engineering approach to study saccadic movements. By using the methods of parameter

estimation and prediction we have produced a model which is remarkably accurate in its tracking and prediction of saccadic eye movements. A detailed account of this model is given in a Ph.D. dissertation by Arun Shah (1977), a copy of which is submitted with this report.

ALPHA PHASE CONTINGENCY ANALYSIS OF THE AVERAGE VISUAL EVOKED POTENTIAL

Although phase analysis of the EEG alpha rhythm has interested a number of neuroscientists, there are, to our knowledge, no published accounts of successful systematic analyses which relate alpha phase contingency to the traditional average visual evoked potential (AVEP). One approach to the analysis of alpha phase contingency effects which has been adopted by a number of investigators (Walter et al., 1946; Walter and Walter, 1949; Turton, 1952; Bekkering and Storm van Leeuwen, 1954; Bechtereva and Usov, 1960; Callaway and Yeager, 1960; Callaway, 1961; Bechtereva and Zontov, 1962; Dustman and Beck, 1965) has been to use what we shall call "phase conditional stimulation techniques" which simplify problems in the collection and management of data by arranging to stimulate only at specified alpha phases. This approach has the advantage of requiring relatively modest instrumentation inasmuch as only one class of phase contingent sample-functions is collected and averaged at a time. Intermediate approaches which provide for the analysis of more than one phase contingency class have been developed by Callaway and Layne (1964) who arranged to stimulate and collect data for 10 different ranges of alpha phase, and by Rémond and Lesèvre (1967) who investigated 4 different ranges of alpha phase. The

disadvantage of these phase conditional stimulation techniques stems primarily from the limitations imposed by the demand for real-time monitoring of the phase variable and for real-time decisions about delivering or withholding stimuli. But possibly of greater importance is the fact that the abandonment of the unconditional or noncontingent pattern of stimulation in favor of the phase conditional pattern also resulted in the production of data which are not easily and directly compared with traditional AVER data.

In contrast with these approaches, we (Anliker and Floyd, 1977a) have elected to expand the traditional AVEP analysis by implementing a completely off-line phase contingency analysis of the sample-functions from which an AVEP can be extracted. In this way, we gain certain fairly obvious analytical advantages from the use of relatively powerful, but time-consuming, digital filtering and classification algorithms. And we obtain results which can be directly compared with the AVEP. Nevertheless, we should point out that our approach has one major drawback, viz., it entails the acceptance of a relatively large computational burden.

EXPERIMENTAL PROCEDURE

The EEG data were collected from 2 healthy normal young men who exhibited prominent alpha rhythms. Electrode couplings were "monopolar," the active electrode being placed over the occipital region of the scalp (3 cm from the midline and 2 cm above the inion) and the inactive reference derived

from yoked electrodes clipped to the right and left ear lobes. The ground electrode was placed over the mastoid process. Right and left occipital responses were recorded simultaneously through amplifiers in a Grass model 78 polygraph onto separate channels of a Honeywell 7600 instrumentation tape recorder. A Grass PS-2 photostimulator, located 12 in. in front of the nasion, was used to deliver flashes (10 msec duration at intensity 4) through the subject's closed eyelids. The flash tube was housed in a sound-attenuating chamber, and white noise was used to mask any flash-tube sounds that might have escaped the attenuator. In addition, the subject was kept both alert and distracted from the flash stimuli by being engaged in a two-way conversation with a laboratory assistant throughout the experimental session.

Digital data processing in this study was carried out on an SEL-840 computer. Prior to the experiment, the computer was programmed to generate a sequence of intervals with values between 2 sec and 8 sec selected from a table of random permutations. This variation in interval size was included to minimize the influence of expectancy which might develop from fixed intervals of stimulus presentation and to counteract the possibility that some lingering stimulus-locked components in the EEG might contaminate the phase-trigger time relationship in the stimulus-OFF control trials. These time intervals were recorded on analog tape along with flash commands (stimulus-ON) and trigger commands (stimulus-OFF)

in separate tape channels. The beginning of each interval had associated with it either a stimulus command or a trigger command. Flash and no-flash trials (i.e., intervals) occurred in simple alternation. This tape was then used to control the stimulus generator during the experiments and to provide signals representing experimental parameters to be recorded onto a second analog recorder along with the EEG signals. These analog recordings were then digitized at 2-msec intervals and stored on digital tape to await further analysis.

Obviously, if one wishes to study alpha phase, it is of great importance that care be exercised to avoid phase distortion in the processing of the EEG signals. Therefore, the digitized EEG data were processed using a nonrecursive phase-distortionless digital filter. The parameters of this transversal filter were (a) bandpass of 6.0 to 18.0 Hz; (b) length of 151 samples at 2 msec per sample; and (c) Hamming weighting. The delay added to the EEG data from the action of this noncausal filter was also introduced into the stimulus channel and the trigger channel to preserve the original time relationship between stimulus commands (or trigger commands) and the EEG time functions.

In subsequent discussions we shall refer to these digital data in terms of sample-epochs or sample-functions corresponding to the various trials. The phrase "stimulus-ON sample-function" will refer to a sample-epoch that begins with a stimulus (flash) and is sampled regularly at 2-msec intervals.

The phrase "stimulus-OFF sample-function" will refer to a similar epoch that begins with a trigger only (indicating that the stimulus was "O F," i.e., omitted).

In the present work we utilize the quadrature definition of phase as originally introduced by Weaver (1956) and applied by Ein-Gal and Lai (1973) to EEG recordings supplied by Anliker. Briefly, the quadrature procedure is as follows: a narrowband time function can be described in terms of frequency, amplitude, and phase values. If the frequency can be specified (either as a constant or a tracked variable), it is possible through quadrature analysis to describe the instantaneous phase and amplitude of the input signal with reference to coherent sine and cosine reference waves of the same mean frequency as the signal. In our experiments the mean alpha frequency was obtained from autocorrelation analysis of the EEG data, the dominant period of the correlogram being taken as the best estimate of the mean alpha period. Quadrature analysis, using this frequency, compares the variable signal with sine and cosine coherent references of the defined frequency and yields an in-phase component and a quadrature component. These two components are used to estimate the amplitude and the phase of the input sample-functions. In the present application, we have made the further specification that the positive peak of the cosine reference wave always coincides with the beginning (time zero) of each sample-function, i.e., coincides

with the moment at which either a stimulus or a trigger-only control is activated. The contingent phase of a particular sample-function is defined in terms of the phase contingency range, of which there are 20 equal 18^{o} ranges, within which the quadrature value of the sample-function falls at time zero. Based upon this quadrature method for the classification of alpha phase, we have developed and applied two forms of phase contingent signal analysis which we shall call (1) phase contingent averaging and (2) phase contingent expansion.

RESULTS AND DISCUSSION

It is instructive to look first at the stimulus-OFF (i.e., trigger-only) phase contingent averages in Figs. P1 and P2. It will be observed that the phase contingent averages are all quite substantial nonzero functions. They are, furthermore, all quite similar except for the phase differences at time zero. The peaks of the amplitude envelopes of these phase contingent averages are all located at time zero. This is to be expected because these phase contingent averages are equivalent to the triggered averages of sato et al. (1962) and the triggered correlations of Boer and Kuyper (1968). Consequently, these phase contingent averages partake of certain basic features of the autocorrelogram. For instance, there is a characteristic fall-off in the amplitude with increasing time distance from time zero. This is accounted for in our phase-incoherence

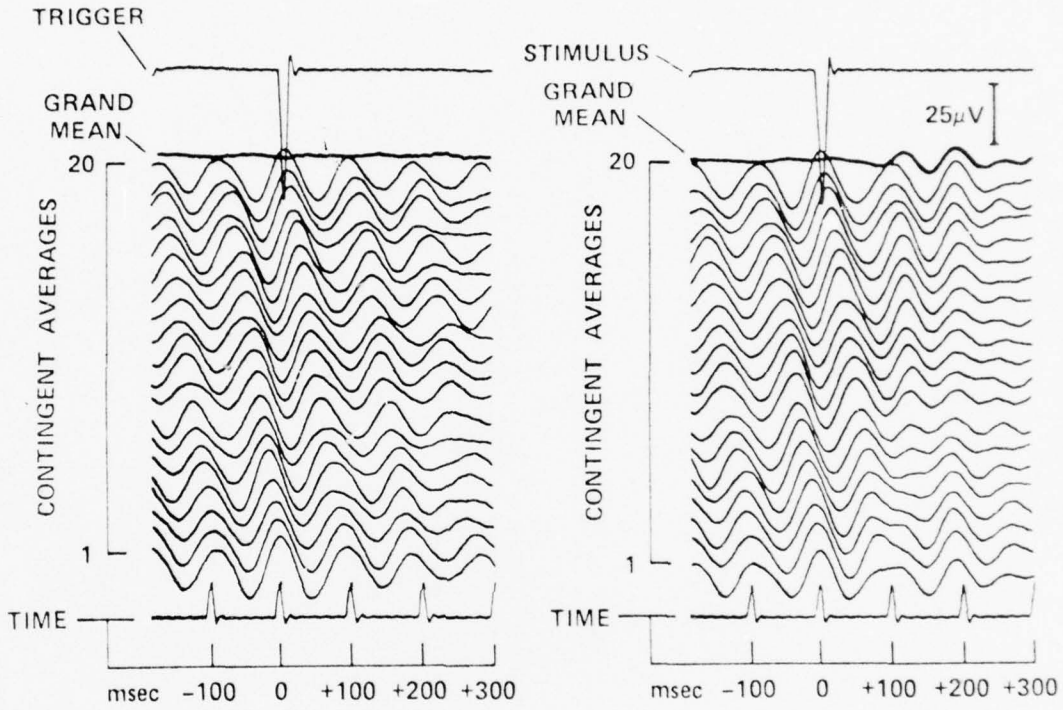

FIGURE P.1.

Phase contingent expansions of AVEP data from subject K.L. The expansion of stimulus-OFF (trigger but no stimulus) data is on the left side of the figure; the stimulus-ON expansion, on the right side. The various contingent averages are stacked as a vertical series, with contingency 1 on the bottom and contingency 20 at the top. The Grand Mean, or mean of the contingent means, equivalent to the noncontingent AVEP, is located immediately above contingency 20 and is indicated by a heavier line. The stimulus (or trigger) time is shown at the top. The average number of sample-functions in each contingent average is 40.

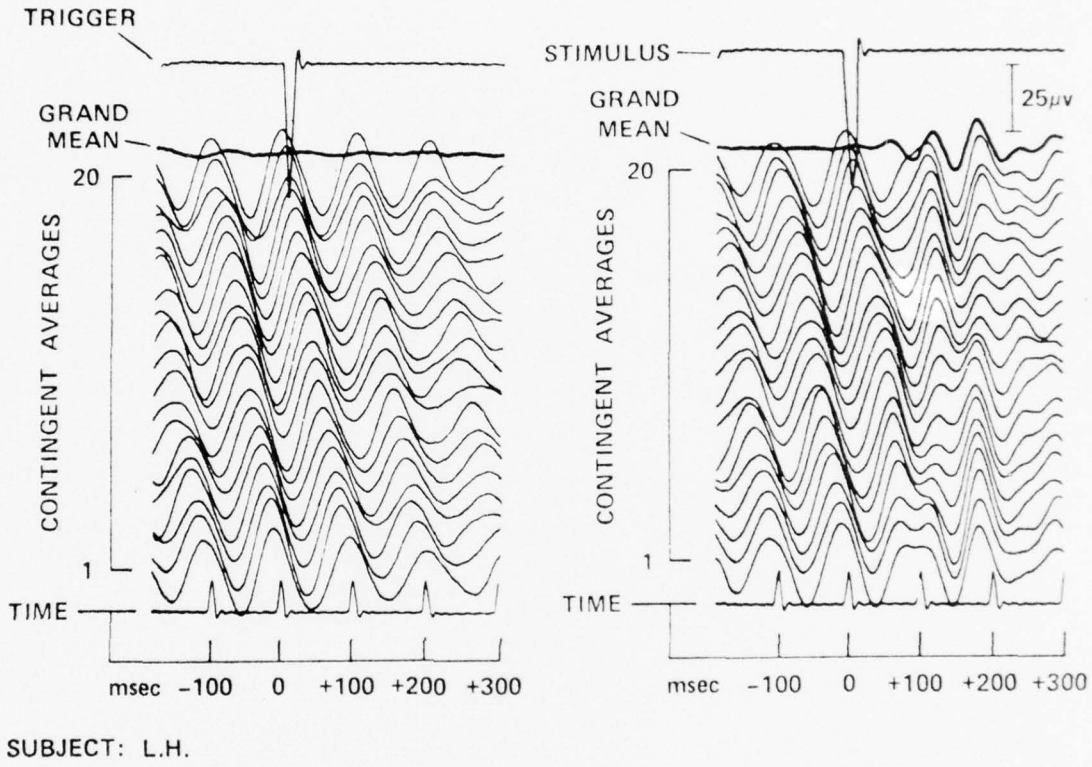

FIGURE P.2.

Phase contingent expansions of AVEP data from subject L.H. The expansion of stimulus-OFF (trigger but no stimulus) data is on the left side of the figure; the stimulus-ON expansion, on the right side. The various contingent averages are stacked as a vertical series beginning with contingency 1 on the bottom. The Grand Mean, or mean of the contingent means, equivalent to the noncontingent AVEP, is located immediately above contingency 20 and is indicated by a heavier line. The stimulus (or trigger) time is shown at the top. The average number of sample-functions in each contingent average is 40.

model (Floyd et al. 1973) by the propagation of phase incoherence as a function of increasing time distance from the classification point which is located at time zero. For a coherent wave train there is no fall-off in the amplitude of the envelope of phase-triggered time averages (or of the envelope of the autocorrelation function). Alpha may be described as a quasi-coherent rhythm, which means, according to the phase-incoherence model, that the rhythm has a coherent mean frequency and a limited amount of phase variance or phase incoherence; the greater the observed phase variance, the faster the rate of fall-off in the amplitude envelope of the phase contingent average. It will be observed in these phase contingent, stimulus-OFF averages that there is sufficient coherence to "predict" some rather definite phase values at those latencies in which the major evoked deflections appear in the AVER (which is equivalent to the Grand Mean of the phase contingent averages). When these phase "predictions" in the stimulus-OFF contingent averages are compared with the phase properties exhibited by their corresponding stimulus-ON contingent averages, there are some striking changes. To facilitate visual examination of these rather complex contingent expansions, we have provided, in Figs. P3 and P4, the same data as in Figs. P1 and P2; some schematic lines have been added to call attention to those features that appear to us to be moved as a result of stimulation. We have used dotted lines to

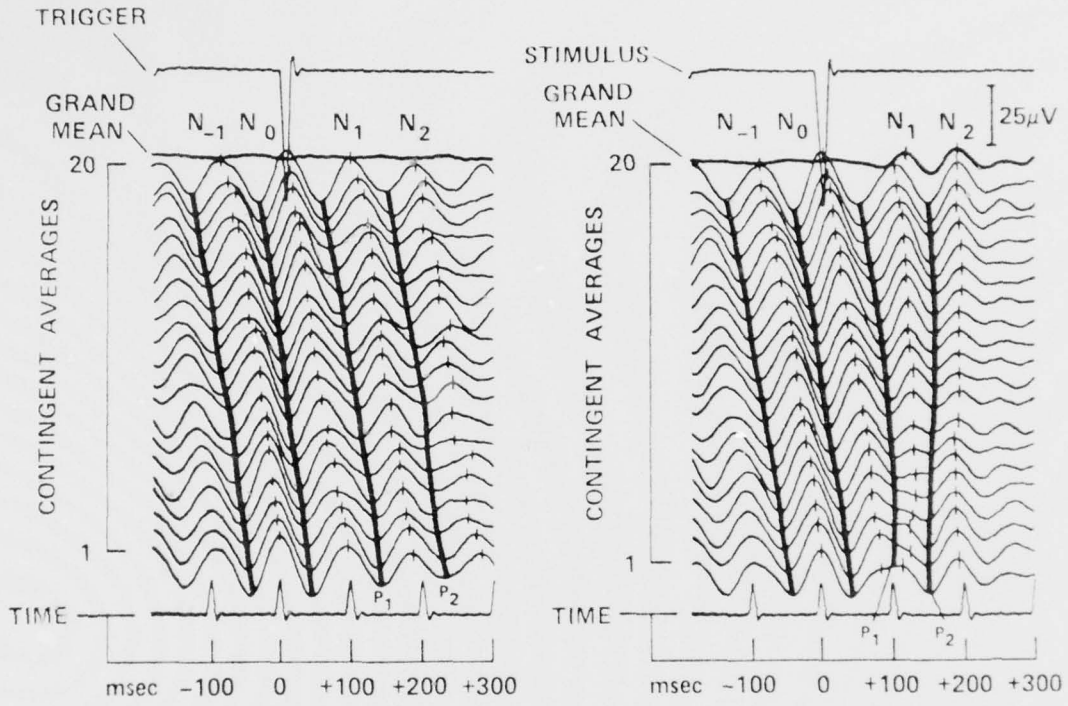

SUBJECT: K.L.

FIGURE P.3.

Schematic overlays on the data of subject K.L. (in Fig. P.1). These overlays indicate the principal differences to be observed between the stimulus-OFF and the stimulus-ON conditions. Notice the phase shift of the N1-P2-N2 components from a diagonal orientation in the stimulus-OFF expansion to a vertical orientation in the stimulus-ON expansion, making these components in all the contingencies agree more or less with the Grand Mean.

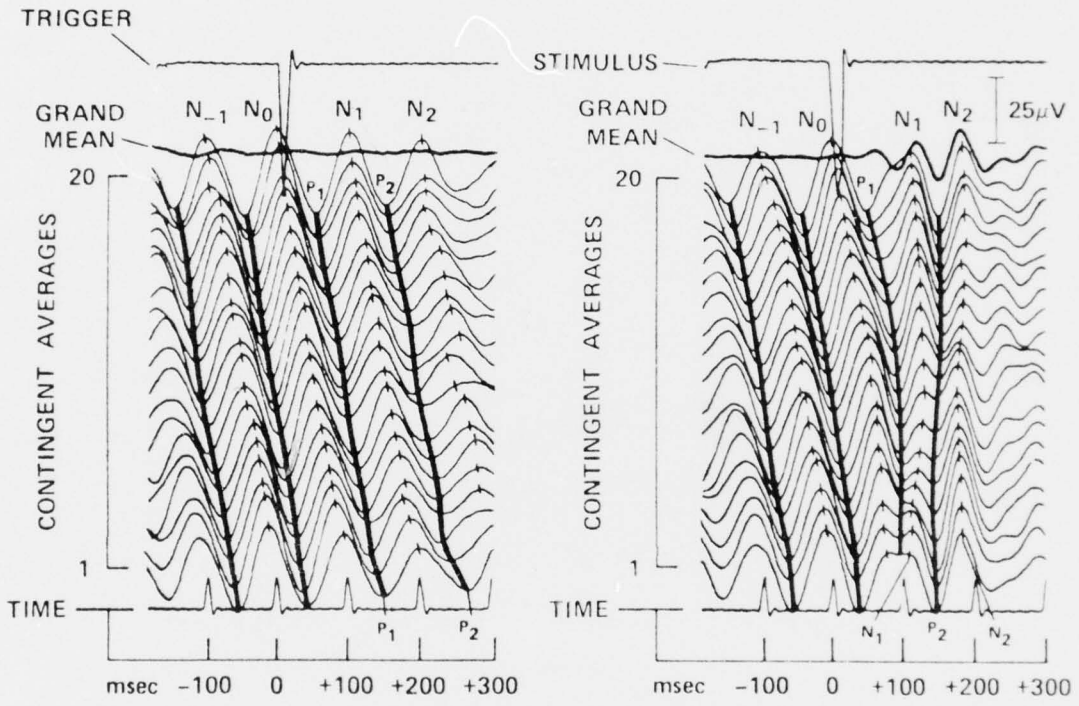

FIGURE P.4.

Schematic overlays on the contingent expansions of subject L.H.(Fig. P.2). These overlays indicate the principal differences to be observed between the stimulus-OFF and the stimulus-ON conditions. Notice the phase shift of the N1-P2 N2 components from a diagonal orientation in the stimulus-OFF expansion (left) to a vertical orientation in the stimulus-ON expansion (right), making these components in all contingencies agree more or less with the Grand Mean. Both subjects show essentially the same effect of the stimulus.

indicate the crests or surface-negative peaks and solid lines to indicate the troughs or surface-positive peaks. In the case of the stimulus-OFF contingent averages, it will be notices that these schematic lines are all essentially parallel diagonals extending upward and to the left from contingency 1 to contingency 20 and connecting corresponding phase points en route. However, when we examine the stimulus-ON phase contingent averages, we notice that these lines are moved in such a manner that the N1, P2, and N2 lines are swung into more or less vertical orientation under the N1, P2, and N2 deflections in the Grand Mean (i.e., AVEP) of the contingent averages. You will notice that the effects are quite similar in both subjects. This surprisingly systematic alteration of the phase properties is interpreted as strong support for the view that the photic stimuli act to produce a phase modulation of the alpha rhythm toward a preferred phase which is consistent with the major N1-P2-N2 deflections in the AVEP. In other words, it is this phase shifting of the alpha from the autonomous-phase locations to the evoked-phase locations which accounts for this portion of the AVEP. That is, the "signal" or evoked component is nothing more than the phase shifting of the autonomous phases to the evoked phases. If this interpretation is correct, there is no need to look for a nonalpha "signal" to account for this portion of the AVEP. In an earlier publication, Floyd et al. (1973) presented data in support of the idea that the so-called "after-discharge" or

"reverberation" which is frequently observed at longer latencies could be accounted for solely on the basis of phase locking which had occurred earlier in the evoked response, i.e., at the time of the appearance of the N1-P2-N2 deflections. In the present work we have produced evidence that such phase alignment does actually occur in this region of the AVEP.

From their study of auditory evoked potentials using amplitude-spectral and phase-spectral methods, Sayers et al. (1974) concluded that "a moderate-level stimulus that is effective in evoking an evident response apparently does not do so by contributing an obvious additive component." And they go on to say that "an effective stimulus acts mainly to phase-control the spontaneous activity." Their view is somewhat similar to ours although they do not identify the spontaneous activity as alpha. Also, whereas our model is based upon a single frequency component (alpha), their model seems to require the existence of a whole series of coherent oscillators (at least 10 harmonics), each with the capacity for being individually adjusted in phase and amplitude. It will be most interesting to see whether the auditory evoked response can be successfully analyzed in terms of phase modulation of a single frequency component such as alpha. The findings of Sayers et al. (1974) certainly suggest that the matter is worthy of careful investigation. It is tempting to speculate that analogous phase encoding mechanisms may operate in all of the cerebral sensory modalities.

AN ALPHA PHASE MODULATION/DEMODULATION THEORY OF CEREBRAL ENCODING/DECODING

In contrast to the widespread assumption that alpha is nothing but noise, we (Anliker and Floyd, 1977a) have proposed a radically different concept which confers upon alpha activities a central role in the encoding and decoding of sensory information. We call our concept the phase modulation theory of cerebral encoding and decoding. According to this theory, alpha serves as a phase coherent carrier on which the stimulus effect is encoded as a phase modulation. That is, the autonomous alpha represents the carrier without modulation while the stimulus-influenced alpha represents the carrier with sensory modulation. The autonomous cortical alpha is considered to represent the cortical-following response to the thalamic pacemaker (the latter being buffered from specific sensory influences). However, the cortical following of the thalamic pacemaker is modulated in the presence of appropriate sensory input to the cortex. The phase modulation theory states that through some process of comparing the pacemaker signal to the response of the corresponding sensory cortex, the monitors (presumably cortical) are able to distinguish between (a) endogenous activities, i.e., those cortical responses which can be accounted for solely on the basis of pacemaker following, and (b) exogenous activities, i.e., those cortical responses which differ from the pacemaker instructions. According to the phase modulation theory, the phase-modulated components

are culled out for further processing while the endogenous or unmodulated phase components are disregarded. In other words, the endogenous activity serves as the undifferentiated medium upon which external influences are encoded as phase shifts or modulations of the autonomous phase. This is regarded as a crucial step in the incorporation of exogenous influences into cerebral processes. One might say that the autonomous activities constitute a low priority input to the cerebral incorporation process and are ordinarily disregarded (i.e., nulled or filtered out) whereas the exogenous activities constitute high priority inputs which will tend to be accepted (passed upward in the neural hierarchy) for higher level processing in attention and memory. How many filters the input must pass before emerging before the "throne of consciousness" is a question which the phase modulation theory does not attempt to answer. But the theory does state that the phase modulation is the sine qua non of cerebral incorporation, the critical step wherein the exogenous influences are translated into a code or format which is the basic language of the cerebrum. This is not to say that frequency and amplitude influences are ignored by the cerebrum (for obviously they are not); however, it does mean that the cerebral "carrier" is pacemaker phase and all time-domain inputs are encoded as phase modulation. Another prediction of the phase modulation theory is that the decoding of the stored phase-modulated signals will require the services of the original carrier frequency; in other words,

playback will require the reactivation of the carrier (i.e., the presence of the autonomous alpha rhythm) for the re-awakening of mnemonic records, i.e., for remembrance. The phase modulation theory predicts that remembrance will be distorted unless the carrier frequency during playback is identical to that during recording. This would account for many of the perceptual distortions which are observed when the autonomous alpha frequency deviates significantly from its normal alert value. Such deviations are observed in normal individuals during states of drowsiness and dreaming, and they can result from the influence of drugs, traumatic brain damage, and metabolic disorders. It should be noted that it is the normal alert alpha that exhibits the greatest phase coherence; when alpha departs from its normal alert frequency there is a reduction in the phase coherence of the autonomous rhythm. Although there are no conclusive EEG records for the alpha frequency differences in psychiatric patients during normal periods of contact and during periods of disorientation, the theory would predict a frequency difference. The low alpha index that characterizes schizophrenics may be a clue that the normal relationship between the pacemaker and the cortex is disturbed; if so, this would tend to cause distortions in the perceptual processes.

It seems to us, however, that some cautionary notes are in order. The alpha phase modulation theory in its present form does not postulate a single alpha pacemaker

for the entire cortex. Rather, we postulate that each portion of the cortex has its own thalamic pacemaker reference along lines suggested by Andersen and Andersson (1968). We would also point out that our model in its present form is concerned primarily with the time domain; we have not at this time attempted to reconcile our theory with an essentially spatial model such as the holographic model of cerebral encoding proposed by Pribram (1971) except to observe that the quasi-coherent alpha activity could provide the coherent reference which is required by a holographic system. Finally, we believe that a careful distinction must be made between the use of the term "phase" in referring to the encoding of the sensory data of visual space (see Pollen et al. 1971), and our use of "phase" as a parameter of the time domain. In contrast, having already utilized the time domain to account for the encoding of spatial phase information (Pollen et al. 1971), Pollen and Trachtenberg (1972) are attracted to the alpha-as-noise hypothesis. They state that "even if it were assumed that all cells in the visual cortex were rhythmically excited by alpha activity and in phase with each other, simple cell responses to visual stimuli in different parts of the receptive field would arrive at the complex cells with different latencies and thus at different periods of the alpha cycle. The alpha activity would facilitate some inputs and inhibit others in a nonpredictable way--a noise process. Thus alpha block, whether by a central active inhibitory mechanism or by a

more general desynchronization of rhythmic activity, might serve to reduce a neural noise level." We believe that our phase modulation theory provides a viable alternative explanation to that advanced by Pollen and Trachtenberg (1972) for the alpha blockage observed during visual attention to nonuniform surfaces. That is, the desynchronization of the cortical response may represent the complex phase modulation of a large cortical area in response to complex spatial patterns of sensory input. These diverse phase adjustments of the cellular "alpha" activities could be expected to greatly increase the phase variance of the cell pool and thereby reduce the magnitude of the potentials measured at a scalp electrode. By contrast, in the absence of sensory phase modulation, a much greater phase coherence in the cortical activity is expected because of coherence of the thalamic pacemaker.

We have concluded, therefore, that simple time ageraging of sample functions of EEG data confounds the influence of alpha phase contingency in the generation of the visual evoked response. By classifying each sample-function according to the phase of the alpha rhythm at the moment of stimulus delivery, it is possible to regroup the set of sample-functions into subsets representing various phase contingency ranges. The time averages of these subsets constitute a phase contingent expansion of AVEP. The results of this analysis can be interpreted as showing that the visual stimulus (flash) acts to phase-modulate the autonomous alpha rhythm and that

this effect accounts for the principal N1-P2-N2 component of the AVEP.

ALPHA PHASE PROBABILITY ANALYSIS OF THE AVERAGE VISUAL EVOKED RESPONSE

The present study was designed to extend the alpha phase contingency analysis and to produce a more quantitative evaluation of the phase contingency effects by reanalyzing AVER data from the standpoint of phase contingent probabilities. We believe that the alpha phase contingency classificatory methods described here represent a definite advance in the techniques available for the analysis of evoked responses and for the quantitative characterization of spontaneous alpha waves.

Phase-Distortionless Filtering

In the analysis of alpha phase it is a matter of utmost importance to take suitable precautions to avoid distorting phase in the process of filtering the EGG signals.

Therefore, in the present study the digitized EEG data were processed using a nonrecursive phase-distortionless filter. The parameters of this transversal filter were (a) bandpass of 6.0 to 18.0 Hz; (b) length of 151 samples at 2 msec per sample; and (c) Hamming weighting. The delay added to the EEG data from the action of the noncausal filter was also introduced into the stimulus record and into the trigger record so as to preserve the original time relationship between stimulus-commands (or trigger-commands) and the EEG sample functions.

DATA-PROCESSING METHODS

Quadrature Analysis

In this study we have employed the quadrature definition of phase which can be described briefly as follows: a narrowband time function can be analyzed into frequency, amplitude, and phase variables. If the frequency can be specified (either as a constant or as a trackable variable), it is possible, through quadrature analysis, to describe the instantaneous phase and amplitude of the input signal with reference to coherent sine and cosine waves of the same mean frequency as the signal. In our study, the mean alpha frequency was obtained by autocorrelation analysis of the stimulus-OFF sample functions of EEG, the dominant period of this autocorrelogram being taken as the best estimate of the mean alpha period. Quadrature analysis, using this frequency, compares the variable input signal with sine and

cosine references of the estimated frequency and yields an in-phase and a quadrature component. These two components are used to estimate the amplitude and the phase of the input sample functions. In the present application, we have made the further specification that the positive peak of the cosine reference wave always concides with the beginning (time zero) of each sample function, i.e., the time at which either a stimulus or a trigger is activated. The contingent phase of each sample function is defined in terms of the phase contingency range (we recognized twenty equal 18^{o} phase ranges) within which the quadrature phase of the sample function falls at time zero. A quadrature phase time function describes the phase difference between the phase of the coherent reference wave and the phase of the EEG sample function at each post-stimulus (or post-trigger) latency.

A quadrature-phase time function is obtained for each EEG sample function. This phase function is used as the basis for further analyses of alpha phase properties.

Phase Probability Analysis

The reader, no doubt, is generally familiar with the procedures for constructing frequency (of events) histograms and for converting these histograms into phase probability distributions. The GRAND PHASE PROBABILITY MATRIX contained in each of the four figures (part E) is really nothing more than a series of phase probability distributions in which each of the probability distributions describes a set of quadrature phase values observed at a particular latency. In

the absence of stimulation one would expect that a sufficiently large population of quadrature phase functions would generate a phase probability matrix with an approximately flat surface because all phases of alpha should be equally probable at every latency. Any significant and consistent deformation of this surface would be considered as evidence for the existence of some sort of phase constraint.

Inasmuch as the grand phase probability matrix provides no account of phase contingency effects, it may prove useful to construct phase contingent phase probability matrices as well. Since our phase classification scheme recognizes 20 different phase contingency ranges, it is feasible to produce a separate phase probability matrix for each of these contingencies. We have, in fact, carried out this analysis. What is the general shape of the phase contingent phase probability matrix? In each case we should expect to obtain a delta function (Dirac) at time zero because all of the samples would be within the same phase range at the time of classification, by definition. With increasing time distance (latency) from the classification point, the propagation of phase incoherence would cause the probability distributions at longer latencies to be less peaked than at earlier latencies because phase variance increases as a function of latency magnitude. On the other hand, in the absence of stimulus-induced phase effects, this "fanning out" of the phase contingent phase probability matrix should be orderly and symmetrical. Consequently, any consistent

and significant differences observed between the corresponding stimulus-ON and stimulus-OFF phase contingent phase probability matrices could be interpreted as evidence for the existence of a stimulus-induced phase probability change.

The grand phase probability matrices in Figs.C1-4 are displayed in pseudo three-dimensional form. While this is attractive to look at, this display distorts probability along the latency axis. To overcome this limitation, two additional plots are provided: modal phase (part C) which shows the location of the probability peaks when projected onto the phase-latency surface; peak amplitude (part B) which shows the amplitude (relative) of the probability peaks above the average probability surface.

It will be helpful to review briefly the manner in which the COMPOSITE PLOT (part D) of the model (or peak) phases of twenty (a complete set) phase contingent phase probability matrices. Since it is not practical to display 80 phase contingent phase probability plots, these data are summarized by extracting the model phase function (i.e., the peak phase at each latency) from each of the matrices and entering these points into a common phase-latency display. By comparing the composite plot with its associated grand phase probability matrix it is possible to discern how the various contingent functions respond to the stimulus. These effects will be mentioned in the Results section.

RESULTS

Our analyses of the experimental data are displayed in Figs. C1-4. These figures are designed to help the reader make his own comparisons of the various analyses contained within each figure and also to aid his comparisons of the corresponding portions of the different figures. Obviously, there are many interesting aspects of these data. Yet, it will be necessary to confine our attention to the most salient features of these data.

Figures C1 and C2 contain the stimulus-OFF and stimulus-ON data of subject K, L. The AVER in Fig. C1a is close to zero amplitude throughout, which is to be expected of stimulus-OFF data. Notice that all of the other plots are consistent with this. The peak amplitude plot (B) shows that the largest non-zero value in the grand phase probability matrix is quite small. The composite plot (D) shows that the contingent modes are not significantly phase constrained. However, we see quite a different picture when we examine Fig. C2 (stimulus-ON). There is a substantial N1-P2-N2 response in the AVER (A). In the grand phase probability plot (E) there is a large deformation of the surface of the matrix in the latency domain that corresponds with the evoked response. The magnitude of this response can be appreciated by looking at the peak amplitude plot (B). What is of great interest to us is the convergence of the modal phases of the contingent probability matrices (D) at latencies corresponding to the evoked response (N1-P2-N2). This strongly suggests

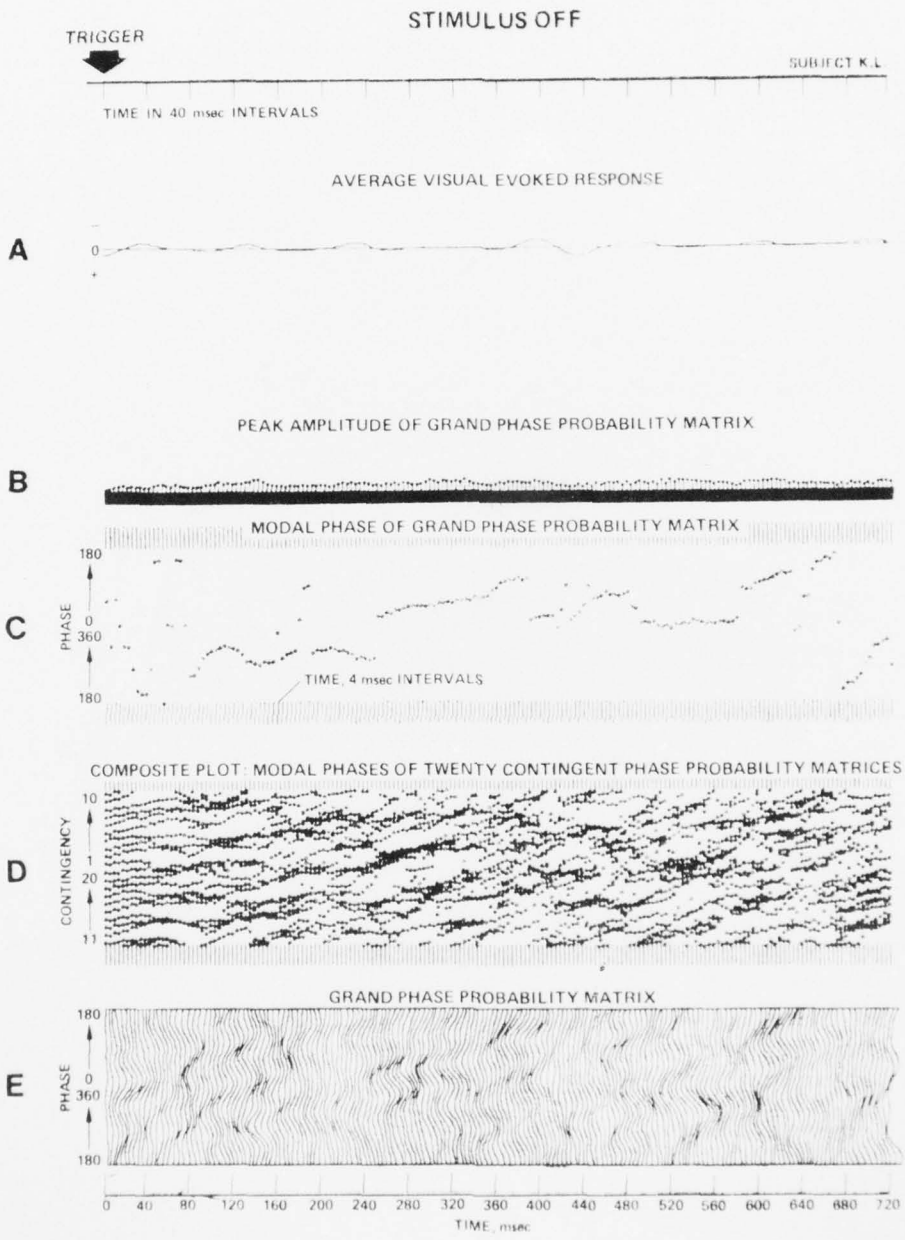

FIGURE C.1. A collocation of the various analyses of subject K.L.'s stimulus-OFF data over a common time-base. Each of the analyses is identified by a capital letter (A - E) located in the left margin and by a descriptive label placed immediately above the graphic material. For example, A = AVERAGE VISUAL EVOKED RESPONSE.

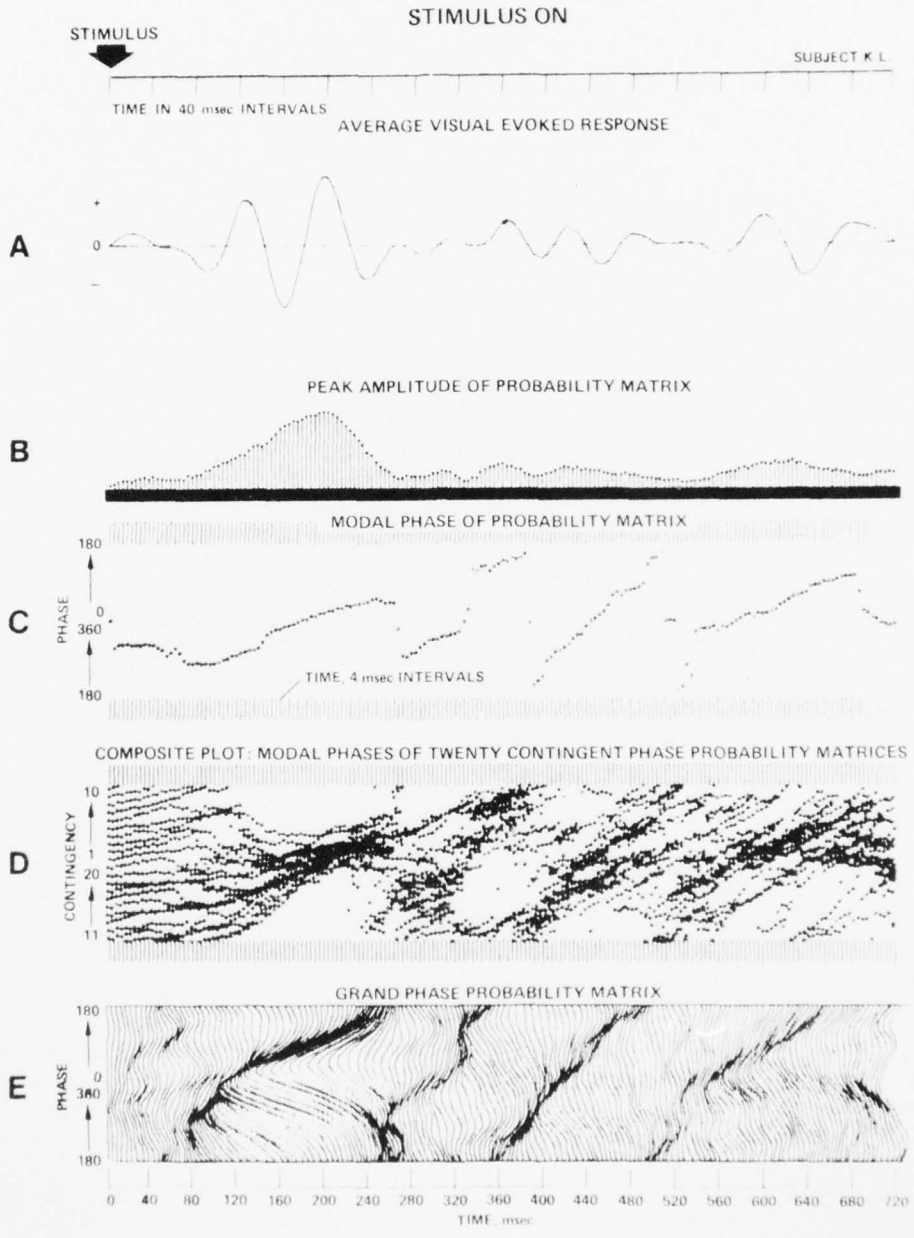

FIGURE C.2. A collocation of the various analyses of subject K.L.'s stimulus-ON data over a common time-base. Each of the analyses is identified by a capital letter (A - E) located in the left margin and by a descriptive label placed immediately above the graphic material. For example, E = GRAND PHASE PROBABILITY MATRIX.

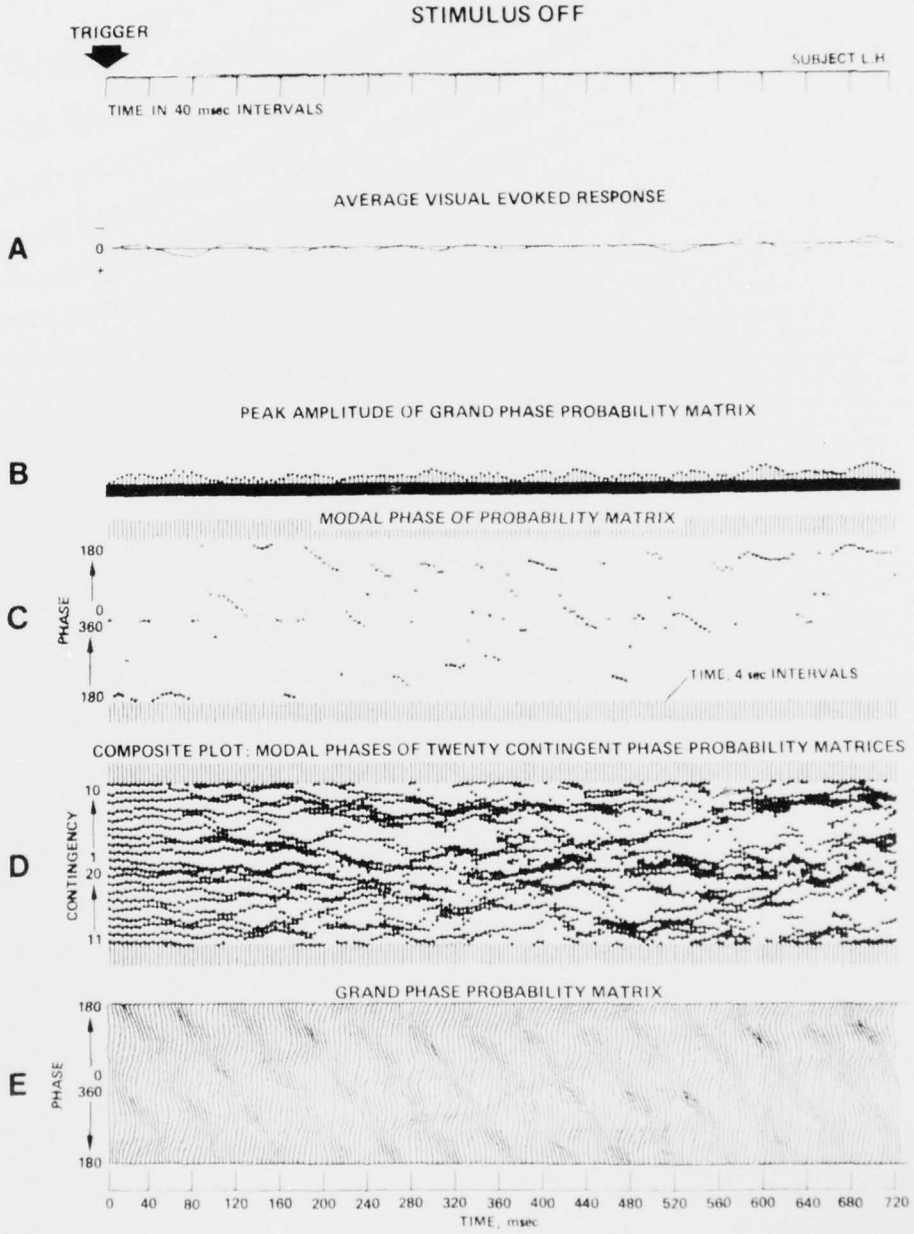

FIGURE C.3. A collocation of the various analyses of subject L.H.'s stimulus-OFF data over a common time-base. Each of the analyses is identified by a capital letter (A - E) located in the left margin and by a descriptive label placed immediately above each type of analysis. For example, C = MODAL PHASE PROBABILITY MATRIX.

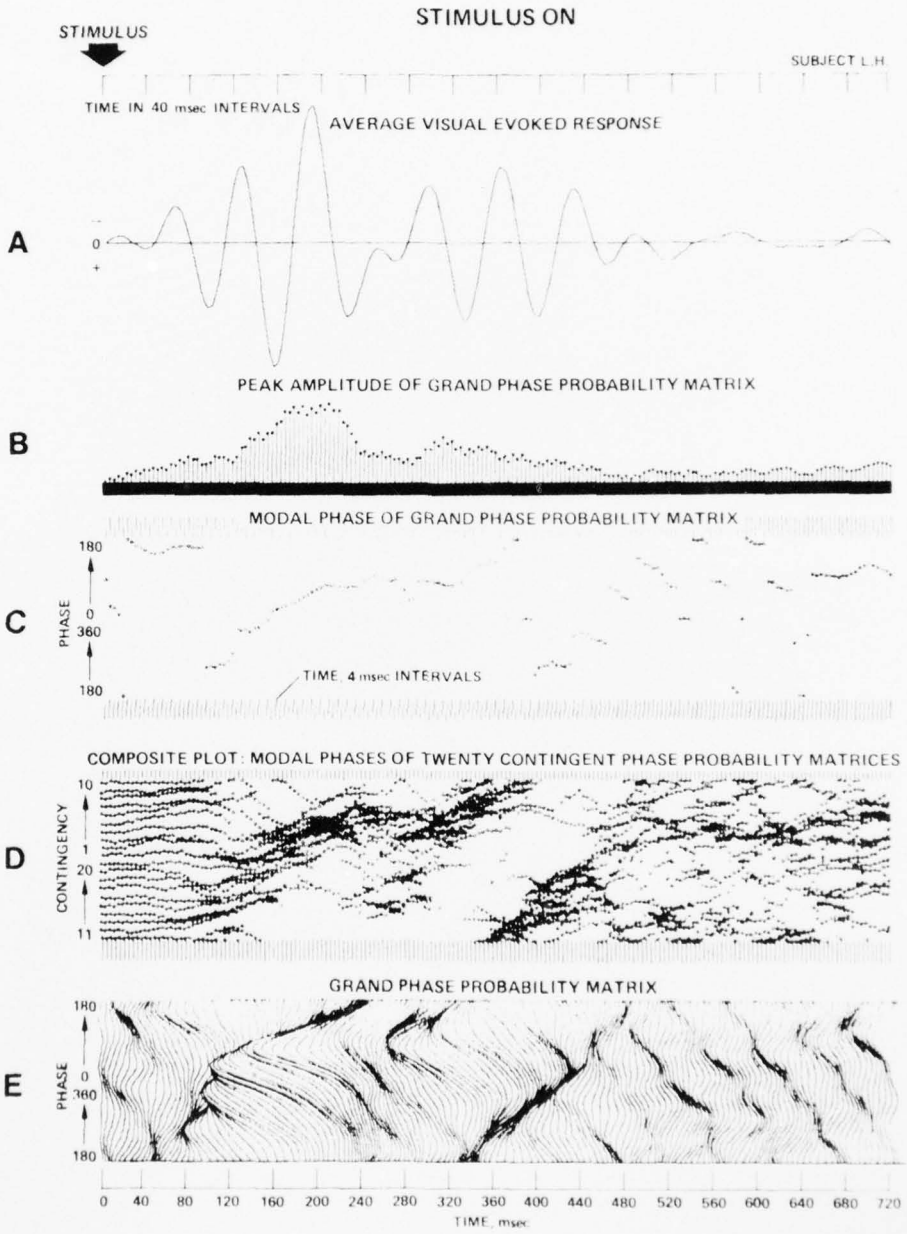

FIGURE C.4. A collocation over a common time-base of the various analyses of subject L.H.'s stimulus-ON data. Each of the analyses is identified by a capital letter (A - E) located in the left margin and by a descriptive label placed immediately above each type of analysis. For example, D = MODAL PHASES OF TWENTY CONTINGENT PHASE PROBABILITY MATRICES.

that the stimulus acts to shift or modulate the autonomous phase toward a preferred phase.

Examination of Figs. C3 and C4 shows that the second subject (L. H.) responded in very much the same manner. That is, both subjects show the convergence of the modal phases (D) and the corresponding deformation of the grand phase probability matrix.

DISCUSSION

It seems to us that the results we have obtained in these two subjects are rather robust. In other words, there is little question that there are substantial phase shifts evident in the stimulus-ON contingent averages; these phase adjustments are imposed upon the autonomous alpha activity, the latter being estimated by the corresponding phase contingent stimulus-OFF averages. It is important to take notice of the fact that all of the phase contingent stimulus-OFF averages are, as should be expected, substantial non-zero averages with a strong alpha-frequency rhythm. For this reason there is a predicted phase for the contingent stimulus-OFF (i.e., autonomous) alpha waves at each latency in the post-trigger time domain). The stimulus appears to shift the alpha phase into conformity with the evoked N1-P2-N2 complex in the AVER. There is some possibility that the various alpha phase contingencies contribute differently to the various latencies; however, this possibility will have to be assessed by further analysis of the data in a subsequent publication. The reader's attention is directed to

the earlier remarks concerning the effects of the propagation of phase incoherence on the shape of the phase probability distributions at the longer latencies; this flattening of the distribution means that the predictability of alpha phase becomes less certain at latencies more distant from the classification point. The alpha generator is only quasi-coherent. Nevertheless, during the latencies occupied by the N1-P2-N2 complex of the AVER, the predictability of autonomous phase contingent alpha phase is still quite good. That is, the phase probability matrix is still exhibiting reasonable peakedness at those latencies. Consequently, the bending of the modal phase probability functions to produce the convergence seen in the stimulus-ON composite plots (part D) shows the ability of stimuli to shift alpha phase toward a "preferred" phase.

In its present form our alpha phase modulation hypothesis does not assume that there is a single alpha pacemaker for the entire cortex. Instead we postulate that each portion of the cortex has its own thalamic pacemaker as suggested by Andersen and Andersson (1968). However, contrary to their model which is based upon barbiturate spindling activity, we assume that the waking brain is characterized by narrow frequency tuning, greater phase coherence, and more coordinated pacemaker activities. We must also point out that our alpha phase modulation model is concerned primarily with the time domain. As yet we have not attempted to reconcile our model with other models which are essentially

spatial in outlook (e.g., the holographic model of cerebral encoding proposed by Pribram (1971) although we are aware that quasi-coherent alpha activity could provide the phase coherent reference which is often required by such models. We would place great emphasis upon the need for making a careful distinction between the use of the term "phase" in referring to the encoding of the sensory data of visual space (e.g., Pollen, Lee and Taylor, 1971) and our use of "phase" as a parameter of the time domain. Having already utilized the time domain to account for the encoding of spatial phase information (Pollen et al. ibid.), Pollen and Trachtenberg (1972) are led to regard alpha as noise. They comment that "even if it were assumed that all cells in the visual cortex were rhythmically excited by alpha activity and in phase with each other, simple cell responses to visual stimuli in different parts of the receptive field would arrive at the complex cells with different latencies and thus at different periods of the alpha cycle. The alpha activity would facilitate some inputs and inhibit others in a non-predictable way--a noise process. Thus alpha block, whether by a central active inhibitory mechanism or by a more general desynchronization of rhythmic activity, might serve to reduce a neural noise level." In contrast, we believe that our alpha phase modulation model provides a viable alternative explanation for the alpha blockage observed during the giving of visual attention to nonuniform surfaces. According to our phase modulation model the desynchronization of the

cortical response under these circumstances may represent the complex phase modulation of a large cortical area in response to complex spatial and temporal patterns of sensory input. These diverse phase adjustments (re-settings?) of the cellular "alpha" activities would greatly increase the phase variance of the cortical cell pool and thereby reduce the magnitude of the potentials under a scalp electrode. The greater alpha amplitude in the absence of patterned visual contact is explained by the greater phase coherence of the cortical cell pool during simple following of the quasi-coherent thalamic pacemaker activities.

SUMMARY

Although simple time averaging of evoked potentials confounds and conceals the influence of alpha phase contingency in the generation of the AVER, it is possible to analyze the role of phase contingency through the use of alpha phase contingent phase probability measurements. The resulting analysis is interpreted as supporting an alpha phase modulation hypothesis of sensory encoding/decoding in the cortex. According to this hypothesis, the sensory stimulus acts to phase-shift the activity of the cortical receptors away from the simple following of thalamic pacemakers. The discrepancy between cortex and pacemaker is regarded as the basis of cerebral encoding of sensory signals.

A NONLINEAR MODEL OF THE HUMAN ELECTROENCEPHALOGRAPHIC SIGNAL DURING PERIODIC PHOTIC STIMULATION

This portion of the Biocybernetics Project is concerned with the development and testing of a mathematical model that is effective in predicting the EEG changes resulting from periodic photic stimulation.

There is little need to justify the development of a mathematical model for characterizing a complex set of empirical behavior. A good model is an essential part of an effective computerized biocybernetic scheme for the prediction and control of actual behavior. If the human brain potentials are to be used in a biocybernetic system as estimators of brain states, it is important to discover the simplest adequate model for characterizing the phenomena of interest. Inefficiency in the model will limit the amount of control achievable.

The human electroencephalogram (EEG) exhibits various complex and poorly-understood responses to repetitive flash stimuli delivered to the eye. In the present investigation, these EEG changes are characterized in terms of the simplest nonlinear input-output model having similar properties of response. Quantitative testing and evaluation of the characterization are achieved through the use of digital simulations of the model's behavior and the use of digital analyses of experimental EEG data. The fit between the model's predictions

and the EEG is good, particularly for periodic stimuli, and extends over a relatively wide range of frequencies. The model is reasonably consistent with what is known about the physiology of alpha wave generation and it is able to unify several apparently disparate phenomena (i.e., they are seen as derivable from the same nonlinear system). The model is also consistent with the sampled-data model (Vossius, 1961; Young and Stark, 1962; etc.) for relationships between saccadic eye-movements, EEG activity, cortical excitability, and visual perception.

As preparation for the development of a model of EEG alpha generation, a large body of relevant literature is examined, including work on the EEG as an autonomous signal, the EEG response to periodic photic stimulation (flicker), the time-averaged response to single stimuli (average visual evoked potentials), thalamo-cortical neurophysiology, thalamic and cortical excitability cycles, visual reaction time, saccadic timing, and short-term visual memory. Of critical importance to this model is the observation that the mechanism of alpha rhythm generation is nonlinear. This and several other key phenomena are noted.

To characterize these phenomena, a simple second-order nonlinear differential equation exhibiting similar behavior is selected. This equation, known as the van de Pol oscillator, is used to model the EEG response to periodic visual stimulation. In this model, the autonomous limit cycle oscillation represents the alpha rhythm, and a coupled

excitation function represents visual field intensity. A perturbation technique for analyzing the response to impulse trains is developed, and numerical verifications are obtained through simulations in a digital computer.

EEG recordings from subjects exposed to sine-modulated or to stroboscopic flashes, on the one hand, and simulated EEG generated by the model, on the other hand, are processed identically and the results are compared. The behavior of the nonlinear model matches the actual EEG responses quite well over a relatively wide range of stimulus frequencies and experimental parameters. Given the low order of the van der Pol oscillator, the number of EEG phenomena which it predicts is rather surprising (and gratifying). However, the fact that the model does not predict transient effects as well as it does steady-state behavior suggests that a higher-order model, such as a network of oscillators, is worth investigating.

It must be recognized that an input-output model of the sort employed here is not a physiological model because there is no plausible correspondence between the model components and particular physiological mechanisms. Much of the success of the oscillator-model used in the present study is due to the fact that it is treated as a canonical model, a tool for classifying phenomena in their simplest form. As such, it provides a quantitative template for any model attempting to account for the same phenomena and it provides a useful set of behavioral classes against which

one may check the performance of an alternative model. Until it has been proved otherwise, we can entertain the hypothesis that the van der Pol type of oscillator is the simplest nonlinear model that exhibits the same classes of behavior as the actual process. In any event, the model provides a useful step in the direction of developing a more complex and more capacious model.

This nonlinear model of the human EEG signal is described in considerable detail in a Ph.D. dissertation by John R. Nickolls (1977), a copy of which is submitted with this report.

REFERENCES

Andersen, P., and Andersson, S.A.: Physiological Basis of the Alpha Rhythm. Appleton-Century-Crofts, New York, 1968.

Anliker, James: "Eye movements: on-line measurement, analysis, and control," a chapter (pp. 185-202) in Monty, R.A., and J. Senders (Eds), Eye Movements and Psychological Processes, Hillsdale: Lawrence Erlbaum Associates, 1976.

Anliker, J.E.: "Biofeedback from the perspectives of cybernetics and systems science," a chapter in Beatty, J., and H. Legewie (Eds), Biofeedback and Behavior: A NATO Symposium, New York: Plenum, 1977 (in press).

Anliker, J.E., and Floyd, R.V.: "Alpha phase contingency analysis of the average visual evoked response," submitted for publication in Electroencephalography and Clinical Neurophysiology, 1977a.

Anliker, J.E., and Floyd, R.V.: "Alpha phase probability analysis of the average visual evoked potential," San Diego Biomedical Symposium, 1977b (in press).

Bechtereva, N.P., and Usov, V.V.:"Technique of intermittent photic stimulation, timed to intrinsic brain potential rhythm, applied to EEG recording," Sechenov Physiol. J. U.S.S.R., 1960, 46: 108-111.

Bechtereva, N.P., and ZONTOV, V.V.: "The relationship between certain forms of potentials and the variations in brain excitability (based on EEG, recorded during photic stimuli triggered by rhythmic brain potentials)," Electroenceph. Clin. Neurophysiol., 1962, 14:320-330.

Bekkering, D.H., and Storm van Leeuwen, W.: "The effect of light flashes triggered by the EEG," Electroenceph. Clin. Neurophysiol., 1954, 6:540.

Boer, E. de and Kuyper, P.: "Triggered correlation," IEEE Trans. Biomed. Eng., 1968, BME-15:169-179.

Callaway, E.: "Day-to-day variability in relationship between electroencephalographic alpha phase and reaction time to visual stimuli," Ann. N.Y. Acad. Sci., 1961, 112: 1183-1186.

Callaway, E., and Layne, R.S.: "Interaction between the visual evoked response and two spontaneous biological rhythms: the EEG alpha cycle and the cardiac arousal cycle," Ann. N.Y. Acad. Sci., 1964, 112:421-431.

Callaway, E., and Yeager, C.L.: "Relationship between reaction time and electroencephalographic alpha phase," Science, 1960, 132:1765-1766.

Clark, M.R., and Stark, L.: "Time optimal behavior of human saccadic eye movement," IEEE Trans. on Automatic Control, June 1975, AC-20, no. 3: 345-348.

Cook, G., and Stark, L.: "Derivation of a model for the human eye-positioning mechanism," Bulletin of Mathematic Biophysics, 1967, 29:153-174.

Cook, G., and Stark, L.: "The human eye-movement mechanism," Arch Ophthal., April 1968, 79:428-436.

Cornsweet, T.N., and Crane, H.D.: "Accurate two-dimensional eye tracker using first and fourth Purkinje images," J. of the Optical Soc. of America, August 1973, 63, no. 8:921-928.

Crane, H.D., and Steele, C.M.: "An accurate three-dimensional eyetracker," Stanford Research Institute: unpublished report.

Dustman, R.E., and Beck, E.C.: "Phase of alpha brain waves, reaction time, and visually evoked potentials," Electroenceph. Clin. Neurophysiol., 1965, 18:433-440.

Ein-Gal, M., and Lai, D.C.: "Error-free representation of EEG signals," Proceedings, IEEE Conference on Systems, Man, and Cybernetics, 1973, 242-243.

Floyd, R.V., Lai, D.C., and Anliker, J.E.: "A model for the photically stimulated electroencephalographic signals," San Diego Biomedical Symposium, 1973, 12:5-16.

Fuchs, A.F.: The Saccadic System from The Control of Eye Movements, P. Bach-Y-Rita and C.C. Collins (Eds), New York: Academic Press, 1971, 343-362.

Morf, M., and Kailath, T.: "A classification of recursive modeling methods," Proceedings of IEEE Conference on Acoustics, Speech, and Signal Processing, Hartford, Connecticut, May 1977.

Nickolls, J.R.: "A nonlinear model of the human EEG signal during photic stimulation," Ph.D. dissertation, Department of Electrical Engineering, Stanford University, Stanford, California, 1977.

Noton, D., and Stark, L.: "Scanpaths in saccadic eye movements while viewing and recognizing patterns," Vision Research, 1971, 11:929-942.

Pollen, D.A., Lee, J.R., and Taylor, J.H.: "How does the striate cortex begin the reconstruction of the visual world?", Science, 1971, 173:74-77

Pollen, D.A., and Trachtenberg, M.C.: "Some problems of occipital alpha block in man," Brain Res., 1972, 41:303-314.

Pribram, K.: Languages of the Brain: Experimental Paradoxes and Principles in Neuropsychology, Englewood Cliffs: Prentice-Hall, 1971.

Rémond, A., and Lesèvre, N.: "Variations in average visual evoked potential as a function of the alpha rhythm phase ("autostimulation")," Electroenceph. Clin. Neurophysiol., 1967, Suppl. 26:42-52.

Robinson, D.A.: "Progress in models of eye movement control," Proc. of the 1972 Int'l. Conference on Cybernetics & Soc., 1972, 19-24.

Robinson, D.A.: "Models of the saccadic eye movement control system," Kybernetik, 1973, 14:71-83.

Sato, K., Honda, N., Mimura, K., Ozaki, T., Teramoto, S., Kitajima, K., and Masuya, S.: "A simplified method for auto-correlation analysis in electroencephalography," Electroenceph. Clin. Neurophysiol., 1962, 14:769-771.

Sayers, B. McA., Beagley, H.A., and Henshall, W.R.: "The mechanism of auditory evoked EEG responses," Nature (Lond.), 1974, 247:481-483.

Shah, A.: "Modeling of saccadic eye movements and EEG alpha rhythm," Ph.D. dissertation, Department of Electrical Engineering, Stanford University, Stanford, California, 1977.

Turton, E.C.: "An electronic trigger used to assist in the EEG diagnosis of epilepsy," Electroenceph. Clin. Neurophysiol., 1952, 4:83-91.

Vossius, G.: "Die regelbewegung des anges," in Aufnahme und Verarbeitung von Nachrichten durch Organismen, ed. VDE. Stuttgart, 1961.

Walter, V.J., and Walter, W.G.: "The central effects of rhythmic sensory stimulation," Electroenceph. Clin. Neurophysiol., 1949, 1:57-86.

Walter, W.G., Dovey, V.J., and Shipton, H.W.: "Analysis of the electrical responses of the human cortex to photic stimulation," Nature (Lond.), 1946, 158:

Weaver, D.K., Jr.: "A third method of generation and detection of single-sideband signals," Proc. IRE, 1956, 44:1703-1705.

Westheimer, G.: "Mechanism of saccadic eye movements," Amer. Med. Assoc. Arch. Ophthalmol., 1959, 52:710-724.

Yarbus, A.: Eye Movement and Vision, New York: Plenum Press, 1967.

Young, L.: "A sampled data model for eye tracking movements," Sc.D. Dissertation, M.I.T., Cambridge, Mass., 1962.

Young, L., and Stark, L.: "A sampled-data model for eye tracking movements," Quart. Progr. Rept. Res. Lab Electr., M.I.T., No. 66, 1962, 370-384.

Zuber, B.L., Stark, L., and Cook, G.: "Microsaccades and the velocity amplitude relationship for saccadic eye movement," Science, 1965, 150:1959-1960.

APPENDIX I

A CLASSIFICATION OF RECURSIVE MODELING METHODS

by

Martin Morf and Thomas Kailath

APPENDIX I

A Classification of Recursive Modeling Methods

I. Introduction

Closely-coupled man-machine systems require efficient estimation and prediction of central nervous system (CNS) activities, such as eye movements and EEG signals. Although the systems responsible for these CNS activities are complex and time-varying, linear systems modeling techniques can be successfully used to predict CNS states. Model parameter estimates can be updated in time, yielding a time-varying linear model. Nonstationarities in observed CNS activities can be modeled with time-varying model parameters, or with a time-invariant linear model with suitable initial conditions.

The flexibility of linear systems modeling has resulted in the rapid development of methods and applications in many areas, particularly seismic data processing and speech modeling. Due to the diversity of these developments, there exists a plethora of methods for estimating the parameters of linear models given input-output data, transfer functions, or covariance functions. Here we present a systematic classification of exact least-squares modeling procedures that are *recursive* (in model-order) and *optimal* in some sense [MKLNV]. Methods which are suboptimal or approximate will only be briefly indicated. Within this framework, we shall point out some new algorithms that have many computational advantages over existing ones. Since we consider state-space, autoregressive moving - average models, and the related ladder realizations, we shall distinguish the following three classes of algorithms: Riccati or square-root type methods, recently developed *fast* algorithms, and their ladder forms. While the first class typically requires computations of $O(n^3)$ or $O(n^2)$ with n equal to the number of model parameters, the *fast* forms only require operations and storage of $O(n)$. Because efficiency is critical in real-time estimation, prediction, and control, these *fast* algorithms are of particular interest here.

In Section II we introduce the modeling problem by reviewing *external* and

internal (linear) models, and consider the types of observed data used to estimate the model parameters. We then introduce a *systematic classification in tableau form* of the various methods to be discussed. It should be stressed here that these least-squares modeling methods can in general only determine the unique *innovations representation* (IR) model [K-S74]. This is usually a finite-dimensional linear model driven by white noise, conceptually the *innovations* or one-step prediction errors. The output of the model represents the process of interest, e.g., eye-movement signals or EEG signals. There exist many such models, but the IR model is the only one which is both stable and stably invertible. The inverse process yields the innovations when driven by the observed data. This means that even though the driving inputs to the IR model are usually not known (e.g., neurological control signals), they can be estimated from the current model parameter estimates. Thus the parameters of the IR model are chosen to produce behavior *statistically* equivalent to the observed data.

In Section III we consider *batch* methods that are best suited for cases where the observed data is accessed in *blocks*. This typically occurs when the data is in the form of a covariance function or system transfer function. A traditional block method based on the Yule-Walker equation [Par] has been used to predict the EEG alpha wave, as indicated elsewhere in this report. Both autoregressive (AR) and autoregressive - moving average (ARMA) models are treated, and the *fast* versions which utilize certain matrix properties are indicated.

In Section IV *recursive* (in time) modeling methods are discussed. These procedures access the input-o[illegible] *sequentially*, and are known as *on-line* methods in the control context [...E], [MKL]. They are ideally suited for real-time biocybernetic applications. Again, the *fast* versions of these algorithms are pointed out.

In section V the new *ladder* (or lattice) - type realizations of the *fast* algorithms discussed in Section III and IV are introduced [IS], [Mo]. These new methods have nice stability properties and good numerical behavior. They also have lowest computational complexity and minimal storage requirements.

In Appendix A we show how to obtain *partial realizations* of the joint impulse response - matching and covariance - matching type, that are both *stable* and *minimal*. In Appendix B, we present an example of our new exact *least-squares recursions* for *ladder-form* realizations, called the "pre-windowing" case [MDKV].

II. The Modeling Problem

The many different types of linear models can be classified into "external" or input/output descriptions, and "internal" or state-space descriptions. We will first consider "external" descriptions, which are sometimes referred to as transfer function type models.

Let us consider a finite-dimensional linear system (FDLS) with inputs $\{ u(\cdot) \}$ and outputs $\{ y(\cdot) \}$. The outputs represent sampled data, such as speech where y's are scalars, or seismic signals from a geophone array where y's are r-vectors. The input-output relationship can be described by an *autoregressive - moving average* (ARMA) model

$$y_i + a_1 y_{i-1} + \ldots + a_n y_{i-n} = w_i \,, \tag{2.1a}$$

$$w_i = b_0 u_i + \ldots + b_q u_{i-q} \,, \quad i \geq 0 \,, \quad n > q \geq 0 \,, \tag{2.1b}$$

where $\{ w(\cdot) \}$ is a moving average of a white noise sequence $\{ u(\cdot) \}$ and the values $\{ y_{-1}, \ldots, y_{-n} \}$ and $\{ u_{-1}, \ldots, u_{-q} \}$ are *initial conditions*. The modeling problem here is to determine the model parameters a_i and b_i. Applying the z-transform

$$y(z) = \sum_{i=0}^{\infty} y_i z^{-i} \tag{2.2}$$

to equation (2.1) yields

$$a(z)y(z) = b(z)u(z) + \{ \textit{terms involving the initial conditions} \} \,, \tag{2.3a}$$

$$a(z) = z^n + a_1 z^{n-1} + \ldots + a_n \,, \tag{2.3b}$$

$$b(z) = b_0 z^n + b_1 z^{n-1} + \ldots + b_q z^{n-q} \,. \tag{2.3c}$$

With zero initial conditions and scalar processes, the ratio of $y(z)$ and $u(z)$ gives the transfer function $T(z) = b(z)/a(z)$. When $b(z) = z^k$, $k \geq 0$ then $\{ w(\cdot) \}$ is a white noise sequence and $\{ y(\cdot) \}$ is called an *autoregressive* (AR) process; the model is referred to as *all-pole*. Alternatively, when $a(z) = z^n$, $\{ y(\cdot) \}$ is a *moving average* (MA) process and an *all-zero* model is obtained.

Turning to "internal" models of the FDLS, we can consider $\{ y(\cdot) \}$ as being generated from $\{ u(\cdot) \}$ with a suitable initial condition $\{x_0\}$ by the *state-space* model

$$\begin{aligned} x_{i+1} &= \Phi x_i + \Gamma u_{i+1} , \\ y_i &= H x_i , \quad i \geq 0 . \end{aligned} \tag{2.4}$$

This model can be chosen to have the given transfer function

$$T(z) = H (zI - \Phi)^{-1} z \Gamma = \frac{b(z)}{a(z)} . \tag{2.5}$$

Note that for convenience we have used a model driven by u_{i+1}, rather than the more commonly used model driven by u_i [MKD] since they can be related [Mo].

Given the transfer function, a simple way to choose the matrices $\{ H, \Phi, \Gamma \}$ is the "observer canonical" form

$$\Phi = Z - a_{1:n} H , \quad H = e_1^T , \quad \Gamma = [b_q^T, 0^T]^T , \tag{2.6}$$

where ^{T} denotes transpose, $a_{1:n} \hat{=} [a_1, \ldots, a_n]^T$, $b_q \hat{=} [b_0, \ldots, b_q]^T$, Z is the "delay matrix": $Z_{i,j} \hat{=}$ *if* $(j-i = 1)$ *then* 1 *else* 0 , and $e_1 \hat{=} [1, 0, \ldots, 0]^T$ is the first unit vector. The state-space model provides a convenient way of computing the covariance function of the output process. Even though the underlying model $\{ H, \Phi, \Gamma \}$ or $\{ a_{1:n}, b_q \}$ is time-invariant, the output process $\{ y(\cdot) \}$ is in general not stationary due to "transients" caused by the initial conditions. However, if Φ is a stability matrix where all eigenvalues have magnitude less than one, then as $i \to \infty$ the transients eventually die out and the process becomes stationary. In the stationary case, the covariance is a function of $|i-j|$ given by $R_y(i,j) = H \Phi^{|i-j|} \bar{\Pi} H^T$, where $\bar{\Pi}$ is the state covariance matrix as $i \to \infty$ (see [DKM]).

ARMA models and state-space descriptions are just two different methods of representing the input-output relations of a FDLS, and they can be closely related to each other using matrix fraction descriptions (MFD's) [DKM]. A lesser-known class of FDLS representations are the *ladder realizations*, which are discussed in section V and are also related to the ARMA and state-space models.

In modeling a process as a FDLS driven by white noise, the observed data is usually available in one or more of the following forms:

a) input-output pairs $\{ u(\cdot), y(\cdot) \}$;

b) impulse response or related sequences such as moments, (or moment estimates, e.g. obtained from input-output pairs);

c) covariance functions or second order knowledge of inputs and/or outputs, (or estimates, e.g. from the impulse response).

Batch methods are used when data is accessed in blocks, as in b) and c); efficient methods for determining model parameters are recursive in order. *Recursive* (in time) methods are appropriate when data is accessed sequentially, as in a). Model parameters can then be estimated recursively both in order and time. Table 1 illustrates the modeling methods that we will discuss; they are divided first into batch and recursive groups, then by model class: Riccati or square-root methods, "fast" algorithms, and their ladder forms. Within each class the name or code for each method appears along with some pertinent references.

Table I: Classification of Least-Squares Modeling Methods

MODEL	METHOD	Riccati, Square-Root Arb. Linear Equations $O(n^3)$ or $O(n^2)$	Chandrasekhar, Fast Cholesky, Levinson Shift-Invariant, Orthogonal Polynomials $O(n^2)$ or $O(n)$	Ladder type Orthogonality $O(n)$
	BATCH			
AR-type	Normal Equation	GE,GS,HH: [GGMS],[Hou]	(Gen.)Levinson,LWR,Szegö-poly.: [K-S74]	Orthog:[IS],[Wak],[MG],[Mo]
ARMA	AR-Part, Minimal Realization	Hankel-Matrix: [HK],[Si]	Töplitz-M.: [Lanc],[Be],[Mass], [DMK], Appendix A	Lanczos Ladder forms: [Mo], [GrMo], [KKM]
	MA-part, Spectral Factorization	Innov.Rep.: [GK]	Chandrasekhar,Fast-Cholesky: [DKM], [Mo]	Fast-Cholesky-Ladder: [Mo], [GrMo]
	RECURSIVE			
AR-type	Recursive Least-Squares	LS,GLS,RE,SQ: [SLG],[MK]	Gen.LWR,RMH: [MKL],[MDK], [MDKV], [Mo]	Recursive Ladder Forms: [MV], Appendix B
ARMA	Augmented Rec. L.-S.	RML1,2,RE,SQ: [SLG],[MK]	Aug.Gen.LWR: [MKL]	Aug.Rec.Ladder Forms: [MV]
AR(MA)	Aug. Nonsym. Rec. L-S	IV,Nonsym.RE,SQ:[SLG],[MK]	Nonsym.Gen.LWR,RMH: [MKL],[Mo]	Nonsym.Ladder Forms: (?)
ARMA	AR-part, Rec. Minimal Realiz.	Spec.IV,Nos.RE,-SQ:Sec.IV	Nonsym.Gen.LWR,RMH,BM-type: Section IV, [Mo]	Nonsym.Ladder Forms: Sec.IV
	MA-part, Rec. Spectral Fact.	Cholesky,RE,SQ: [GK],Sec.IV	Fast Cholesky,RMH5: Sec.IV,[Mo]	Rec.Feedb.Ladder: Sec.IV,[Mo]

III. Batch Methods

When observed data is available in blocks, *batch* modeling methods are convenient. We will first consider AR models because of their widespread use [Makh]. The recently developed "linear predictive coding" (LPC) speech compression schemes, for example, are a direct application of least-squares fitting of AR models [AH], [IS], [Wak] (for a survey see [MG]). AR models have also been very useful in statistics [Par], [BJ], spectral analysis [UB], [Aka], and multichannel geophysical applications [Rob], [WR].

Normal Equations

It is well known in least-squares problems that the parameters of an AR model satisfy a set of *linear* equations called the *normal equations* (see [K-S74], [MG]):

$$\mathbf{a}_n^T R_n = [1, -a_1, \ldots, -a_n] R_n = [R_\epsilon, 0, \ldots, 0] \quad . \tag{3.1}$$

An alternate form is the Yule-Walker equation [Par]:

$$\mathbf{a}_{1:n}^T R_{n-1} = [a_1, \ldots, a_n] R_{n-1} = [r_1, \ldots, r_n] \quad . \tag{3.2}$$

In both forms R is a covariance matrix and the a_i's are the "predictor" or AR model parameters. R_ϵ is the "prediction error" or innovations covariance, a non-increasing function of n (typically the model order). In speech processing, two popular methods of obtaining the normal equation are the "autocorrelation" method and the "covariance" method [MG], but there exist many ways of estimating the covariance R_n [BJ], [MDKV], [Di]. General methods for solving such linear equations include Gaussian elimination (GE), Cholesky decomposition, Householder transforms [Hou], [GGMS]; however, they all require computations of $O(n^3)$.

The Levinson-Wiggins-Robinson (LWR) Algorithm

An algorithm that requires only $O(n^2)$ computations for the recursive solution of normal equations with *Töplitz* covariance matrices (corresponding to an assumption of *stationarity* of the process) was first described by Levinson [Lev] and later extended by Wiggins and Robinson [WR]. By making use of certain

shift-invariance properties of Töplitz matrices (the i,j-th entries are only a function of $i-j$), this algorithm solves the normal equation via a set of recursions that update the AR parameters or the predictor parameters in increasing order [Rob], [K-S74]. The Levinson recursions are also closely related to the orthogonal Szegö polynomials [Sze], [GS], [KVM]. Levinson's algorithm plays a central role because it can be generalized to handle multi-channel data [WR], multi-dimensional or image processing problems [LKM], nonstationary processes with "*shift-low-rank*" [*Mo*], [*FMKL*], ladder realizations [MV], [MVK] and ARMA or state-space models [MKD], [MKL], [DKM].

ARMA Models

In state-space terms, the problem is to find a triple $\{ \mathbf{H}, \Phi, \Gamma \}$ such that $T_i = \mathbf{H}\Phi^i\Gamma$, where $\{ T_i \}$ is a given set of "first order" data characterizing the impulse response of a linear system. This is the *partial realization* problem [KFA], [DMK]. The central role in this realization theory is played by the *Hankel* matrix with entries $\boldsymbol{H}_{i,j} = T_{i+j-1}$. The columns or rows of this matrix are known to span the state-space, so any method for finding a basis is a viable realization method. Of particular interest are methods for finding the smallest basis resulting in minimal order n realizations [*HK*], [*Si*], [*YT*]; they all require $O(n^3)$ operations.

From a transfer function point-of-view, the partial minimal realization problem is that of finding two relatively prime (matrix) polynomials $a(z)$ and $b(z)$ such that the given power series $T(z)$ matches say k terms of the expansion of $b(z)/a(z)$. This is the classical Padé approximation problem, which in the scalar case yields $T(z)\, a(z) = b(z) + \{ \textit{terms in } z^{-i},\ i > k-n \}$. Equating coefficients of z^{-i}, $0 \le i \le n$, we get

$$\boldsymbol{H}_n \mathbf{a}_n = 0_n \ , \ \text{or} \quad \boldsymbol{H}_n [a_n , \ldots , a_1]^T = - [T_{n+1}, \ldots , T_{2n}]^T \ , \tag{3.3}$$

where $\bar{\boldsymbol{H}}_n$ is a Töplitz matrix containing the reversed column ordered Hankel matrix $\boldsymbol{H}_n$. Note the similarities here to the normal equation (3.1) and to Prony's method [MG].

Again, standard methods could be used to solve for $\mathbf{a}_n$, but they all require computations of $O(n^3)$. However, if one takes advantage of the structure of the Hankel or Töplitz matrices, fast algorithms can be found. Such algorithms have been developed (in a coding theory context) by Berlekamp [Be] and for minimal realization by Massey [Mass]. Multivariable versions have also been developed [Mo], [DMK]. These recursions are strikingly similar to Levinson's recursions; the Berlekamp-Massey algorithm can also be related to orthogonal Lanczos polynomials [Lanc], [Mo]. An alternative method for obtaining *stable partial minimal realizations* is discussed in Appendix A. It can also be derived by considering a Gram-Schmidt (GS) orthogonalization on the columns of the Hankel matrix $\boldsymbol{H}_n$ or the Töplitz matrix $\boldsymbol{H}_n$ [Mo], or more generally by using projection methods [KKM].

Spectral Factorization and Innovations Representations

The problem of obtaining a model of a process $\{ y(\cdot) \}$, given its covariance function or second order information, is called *stochastic realization*. We are interested in representing $\{ y(\cdot) \}$ as the output of a linear model driven by white noise. In general, there exist many such models, however the only stable and stably invertible model is the (unique) *innovations representation* (IR) [K-S74]. The inverse model is the whitening filter that produces a white noise process, the innovations $\{ \epsilon(\cdot) \}$, when driven by the observed data. In discrete-time or time series analysis, the innovations are the one-step prediction errors of the observations. If the process $\{ y(\cdot) \}$ is stationary, the problem of obtaining the IR essentially reduces to one of spectral factorization.

An efficient method for obtaining the spectral factors of $S_y(z)$

$$S_y(z) = b(z)\, b(-z) \,/\, a(z)\, a(-z) \quad , \tag{3.4}$$

is given as a two-step procedure in [DKM], [MKD], [Mo]. In the stationary and scalar process case considered here, the truncated or one-sided power spectrum $S_+(z)$ of $\{ y(\cdot) \}$ is formed from the covariance sequence. A minimal realization algorithm

is then used to approximate $S_+(z)$ by $q(z)/a(z)$, where $1/a(z)$ is the AR-part of the desired model. From

$$\begin{aligned} S_w(z) &= a(z)\, S_y(z)\, a(z^{-1}) = a(z)q(z^{-1}) + q(z)a(z^{-1}) \\ &= b(z)\, b(z^{-1}) \end{aligned} \tag{3.5}$$

it can be seen that the spectral factorization problem for $\{ y(\cdot) \}$ is now reduced to the simpler factorization problem for a MA-process $\{ w(\cdot) \}$, where the factor $b(z)$ is the MA-part of the desired model. It can be obtained by Cholesky and other factorization procedures. Thus $\{ y(\cdot) \}$ is modeled by the cascade of the AR and MA parts driven by the innovations, a white noise.

In the time-domain this corresponds to a factorization of the (stationary) covariance (a Töplitz matrix) into triangular factors. The "fast-Cholesky" factorization given by [Mo] is an efficient algorithm for stationary and non-stationary covariance matrices with "shift-low-rank". It should be noted that many popular covariance estimates have this property.

IV. Recursive "in Time" Methods

When the observed data is available as input-output pairs that must be accessed sequentially, recursive modeling methods are the most appropriate. Many recursive least-squares methods have been developed in the identification and control area; they typically involve solving Riccati-type equations and have computation and storage requirements of $O(n^3)$ and $O(n^2)$ respectively. Recently "fast" algorithms have been developed with reduced computations and storage of $O(n)$ using ARMA or ladder realizations.

An important set of least-squares recursive methods for AR-type models is described in detail in [SLG] and more recently in [MKL]. The discussion includes least-squares (LS), weighted least-squares (WLS), generalized least-squares (GLS), instrumental variable (IV) and recursive maximum-likelihood (RML1,2) methods for ARMA models. All these methods solve a Riccati equation that recursively updates the inverse of the matrix appearing in the normal equation of the problem.

An alternative to the Riccati equations are the square-root forms discussed for instance in [MK]. They make use of the numerically preferrable orthogonal transformations [Hous], [GGMS].

A special case of the IV method is obtained by using the n-step delayed outputs as instrumental variables. This can be shown to be equivalent to a minimal realization problem given (estimated) covariances $\boldsymbol{R}$. Recall that in the given (estimated) covariance case we discussed a two step procedure. The first step was to obtain a minimal realization, or rational approximation of $S_+(z)$ by $q(z)/a(z)$, or in the time-domain of $\boldsymbol{R}_+$ by $\boldsymbol{Q}\,\boldsymbol{A}^{-1}$ [DKM], [MKD]. In matrix notation we obtain

$$\boldsymbol{R}_+\,\boldsymbol{A} = \boldsymbol{Q}\ , \tag{4.1}$$

where $\boldsymbol{A}$ and $\boldsymbol{Q}$ are banded matrices of "band width" n, if the underlying linear model is of that minimal order. The first column of (4.1) corresponds to equation (3.3)

$$\boldsymbol{R}_n\,\boldsymbol{a}_n = \boldsymbol{H}_n\,\boldsymbol{a}_n = 0_n\ . \tag{4.2}$$

where R_n is the (2,1) block entry of the appropriately partitioned triangular Töplitz matrix R_+. The matrix R_n is the cross-covariance of the last n observations and the same set n time-steps delayed. R_n clearly plays here the role of the reversed ordered Hankel matrix H_n. By noting that the Riccati-type equation can be interpreted as a recursion for (low rank) updating of the inverse of a matrix, we see that the IV method can be viewed as a recursive (in time) updating procedure for the minimal realization solution for a in equation (3.3).

Fast Algorithms for Recursive "in Time" Methods

In [MKL] the development of "fast" algorithms for the recursive least-squares methods is discussed in detail. Efficient recursions for time and/or order updates for AR-type models were first derived in [Mo]. The basic idea there was the observation that the matrices encountered in many least-squares problems have a "shift invariance" or a "shift-low-rank" property. It can be characterized by the (low) rank ρ of the shifted difference of a matrix M: $\rho [M - Z^T M Z]$, where the "delay" matrix Z was defined in Section II. This property is generated by the fact that these matrices are sums of products of Töplitz or Hankel matrices. It can also be used to obtain fast Cholesky algorithms for MA processes, thus obtaining *recursive* whitening filters of the AR type (e.g. RMH5 algorithm in [Mo]).

Similarly we can obtain general LWR recursions in order and time for AR processes, i.e. updating the MA prediction (whitening filter) parameters a_i. A surprising feature of the fixed-order recursive-in-time algorithms is that explicit updating of the covariance estimate is unnecessary, basically because the model parameters are an *implicit characterization of the covariance*. Since the details of these algorithms can be found in [MKD],[DKM],[Mo],[FMKL] , we shall only give a comparison of the LWR-type algorithms, assuming that the reader is already familiar with the Levinson recursions as described in [Wie], [WR], or more recently [K-S74].

The recursions for the generalization of the LRW algorithm for covariance matrices exhibiting a *shift invariance* property have a very similar form to the

original Levinson recursions. However, in contrast to the two solutions required in the LWR algorithm, the so-called forward and backward predictors, we require in general more solutions for non-Töplitz matrices. It turns out that the "shift-low-rank" ρ of the covariance matrix, regardless of its size, is equal to the number of solutions required in the recursions. For the case of covariance estimates that can be written as products of two Töplitz matrices (typically containing input-output data), the number of solutions required is at most four in the scalar case, and $2m+2$ for m-channel data.

For combined ARMA models we can either attempt to model first the AR or the MA part and then try to estimate the remaining part of the model. In the batch methods of Section III we discussed ways of obtaining the AR part first via minimal realization and then the MA part via spectral factorization. The other order of first obtaining the MA part could be obtained by working with (an estimate of) the inverse of the covariance matrix, the so-called *information* matrix [MK].

The cascaded approach can be carried out also in time recursive form by estimating the AR part via a (fast) recursive form of the IV method, as discussed above, cascaded by the fast Cholesky recursions for a MA process (e.g. RMH5 in [Mo]). The only difficulty now is that both parts estimate the models in the so-called controller or "tapped delay line" realization, a dual form to the observer form, which cannot be merged by inspection.

The joint innovations representation approach discussed in Section III and Appendix A is ideally suited for recursive in time methods. Even though the driving inputs (conceptually the innovations) are usually not available, they can be replaced with their best estimates obtained by using the best current parameter estimates. This is clearly only possible for methods with sequential data access. It turns out that this seemingly "suboptimal" approach of substitution has itself optimality properties (see e.g. [SLG]); a similar situation occurs in detection of unknown signals, and in the famous separation result of linear quadratic control using state estimates [KFA]. The recursive maximum likelihood methods in [SLG] and [MKL] can be derived from such an approach. Once estimates of current

prediction errors are obtained, they can conceptually be treated as known data, and entered for instance in normal equation expressions. The only problem that might arise lies in theoretical proofs of the convergence of such methods.

V. Ladder Realizations

In Section II we discussed the realization of a given transfer function $T(z) = b(z)/a(z)$ via transfer function or state-space type models such as ARMA, controller, or observer canonical forms. If the roots of $a(z)$ are known, $T(z)$ can be represented by a partial fraction expansion. Using polynomial evaluation formulas we can obtain the so called Jordan canonical or parallel decomposition form (even for matrix transfer functions) [Mo]. The Jordan form has the nice property that stability can be checked by merely inspecting the magnitude of the roots . Finding the roots, however, is numerically sensitive.

The ladder (or lattice) canonical realizations provide a very promising alternative. They also have the property that stability and even minimum-phase can be checked by inspecting the PARCOR or *reflection coefficients* [IS], [Wak], [MG], [Mo], [Cla2]. In contrast to the methods for finding roots of polynomials this requires only a *finite* algorithm, the Schur-Cohn test for stability. This is actually equivalent to the Levinson or orthogonal Szegö polynomial recursions performed in decreasing order on $\mathbf{a}_i$ or $a_i(z)$, [K-S74]. Given the stationary covariance matrix $\boldsymbol{R}$, i.e. second order information, the $\mathbf{a}_i$'s and the reflection coefficients can be computed via the LWR algorithm. From a stochastic process point of view we can identify these coefficients with the *partial correlation (PARCOR) coefficients* or singular values. They also have physical significance in the scattering theory of waves [IS], [Wak], [K-S74], [MV], [MVK].

The Levinson algorithm can actually be carried out using *only* the reflection coefficients as parametrization, since the inner product k_i of a vector $\boldsymbol{r} \hat{=} [r_1, \ldots, r_i]^T$ and the coefficient vector $\mathbf{a}_{i-1}$ can be obtained as the current output k_i of a ladder realization driven by a previous input sequence containing $\{r_1, \ldots, r_i\}$. Similarly we can translate other algorithms for AR models, such as the various generalized Levinson algorithms [Mo], [DKM], [LKM], Appendix A, into their ladder form equivalents as in [MVK], [MV], Appendix B. These forms are of interest by virtue of their stability properties and their numerical robustness -- they typically require sample correlation operations. These forms have also canonical [Mo] and

invariance properties [MVK], as well as minimal storage requirements for modeling algorithms, as seen from a comparison of the the recursions for the PARCOR and the a_i parameters in Appendix B.

ARMA models

Recall that the first step of realizing an ARMA model in Section III was a minimal realization problem. The solution to this MR problem can be carried out in ladder form by using a Berlekamp Massey (BM) - type algorithm. These recursions are actually very similar to the LWR recursions, as noted in [Mo]; therefore we can use an analogous derivation to obtain ladder forms for the BM recursions, as presented in [GrMo].

Alternatively, the joint IR approach explored in Appendix A, leads (even for scalar) processes to the theory of multichannel ladder realizations of the AR type discussed above. Since we embedded the underlying ARMA model in a two channel AR model, the IR model will again be of order $2n$, i.e. non-minimal. This would also hold for a ladder realization.

Minimal models can be obtained by merging the AR and MA parts in the observer form. It is also possible to obtain a minimal rational ladder form [MG], [Mo]. The basic idea in state-space terms is to add a suitable input matrix (Γ) or output matrix (H) to a ladder form realization of the AR part of the model; this is possible since the ladder forms are controllable (or their dual observable) [Mo].

The second step of the stochastic realization procedure of Section III requires a spectral factorization for the determination of the MA part. As indicated, we need to determine the triangular factors of the (banded) covariance matrix of the MA process. They can be obtained from the Cholesky factors or the RHS of the LWR recursions. Similarly it is possible to obtain the ladder realizations of the MA part, since the fast Cholesky recursions "by rows" have the form of a state-space equation with a dynamic matrix Φ that has precisely the same form as the Φ matrix of a feedback ladder form in state-space notation [Mo]. As for the LWR recursions, there similarly exists a ladder form of the fast Cholesky algorithm that requires

only reflection coefficients as parametrization.

Ladder Forms for Recursive "in Time" Methods

The ladder forms for *exact* least-squares solution to AR modeling have been developed in [MV]. In Appendix B, we shall present the simplest one of the many possibilities of such ladder forms, the "prewindowing" case [MDKV]. It is interesting to note that the partial correlation coefficients are computed as sample cross correlations between the "forward" and "backward" prediction errors as expected from the stochastic derivations of the ladder forms [Wak], [Mo], [SKM].

The ladder forms of the GLS and RML1/2 methods discussed in [SLG], [MKL] and Section IV can be obtained by embedding the models in an appropriately augmented AR model as in Appendix A . The IV method led to nonsymmetric Riccati equations, therefore the fast versions also require a nonsymmetric form of the LWR recursions. However it is clear that these recursions are then of the BM type since this algorithm also works with nonsymmetric (though triangular) Töplitz matrices. Therefore, we could obtain "nonsymmetric" ladder forms of the type given in [GrMo]. Although the final algorithms of these ladder forms are simple to implement, the exposure of the "shift-invariance" is in general nontrivial [MV].

Our preliminary experience with the numerical properties of these algorithms has been very encouraging; in general ladder realizations are superior to direct forms for computing estimates of the coefficients of $a(z)$ and $b(z)$. Stochastic approximation or gradient type methods using ladder forms can be obtained easily, e.g. [SV]; however, they have drawbacks similar to other stochastic approximation methods, especially for covariance matrices with extreme eigenvalue distributions.

Conclusions

We have presented a number of modeling methods in a framework which includes the new ladder-form realizations. Of the three types of procedures considered, e.g. Riccati or square-root methods, fast methods utilizing matrix shift-invariance properties, and ladder-form methods, the latter two are most suited to problems requiring efficiency. In particular, the recursive in time versions lend themselves to on-line or real-time applications because the input-output data is accessed sequentially one sample at a time. Also, new parameter estimates are available at each sample time, which facilitates on-line control problems.

The ladder-form methods also have the desirable property that their stability can be checked merely by inspection of the reflection coefficients. In addition, they are numerically robust since the major operations are sample cross correlations. They also have minimal storage requirements for least-squares modeling methods. The structure of the ladder-type realization suggests that the reflection coefficients may have physical significance in the process being modeled. For instance, in speech models, the PARCOR coefficients correspond to acoustic wave reflection coefficients in the vocal tract.

Appendix A: Stable Partial Minimal Realizations

In this appendix we show how to obtain stable partial minimal realizations of the joint impulse-response and covariance-matching type. It will turn out that we can obtain an ARMA model by imbedding it in a two channel AR modeling problem. Given an ARMA model as represented by the difference equation of section II, we can rewrite it as (let $q = n$)

$$y_t + a_1 y_{t-1} + \cdots + a_n y_{t-n} - b_1 u_{t-1} - \cdots - b_n u_{t-n} = b_0 u_t , \tag{A1}$$

or $$\mathbf{a}^T \mathbf{y}_t - \mathbf{b}_1^T \mathbf{u}_t = b_0 u_t , \quad \text{where}$$

$$\mathbf{a}^T = [\, 1 , a_1 , \ldots , a_n \,] , \quad \mathbf{y}_t^T = [\, y_t , \ldots , y_{t-n} \,] ,$$

$$\mathbf{b}_1^T = [\, 0 , b_1 , \ldots , b_n \,] , \quad \mathbf{u}_t^T = [\, u_t , \ldots , y_{t-n} \,] .$$

Now consider the following augmented equation

$$\begin{bmatrix} \mathbf{a}^T & -\mathbf{b}_1^T \\ 0 & e_1^T \end{bmatrix} \begin{bmatrix} \mathbf{y}_t \\ \mathbf{u}_t \end{bmatrix} = \begin{bmatrix} b_0 u_t \\ u_t \end{bmatrix} , \tag{A2}$$

(e_1^T is the first unit vector). This equation can be interpreted as an AR model for the joint process $\{ \mathbf{y}, \mathbf{u} \}$ [Mo], since the RHS is equal to the joint innovations of $\{ \mathbf{y}, \mathbf{u} \}$, since

$$\boldsymbol{\epsilon}_t = \begin{bmatrix} \epsilon^y_t \\ \epsilon^u_t \end{bmatrix} = \begin{bmatrix} y_t - \hat{y}_{t|t-1} \\ u_t - \hat{u}_{t|t-1} \end{bmatrix} = \begin{bmatrix} b_0 u_t \\ u_t \end{bmatrix} . \tag{A3}$$

Deterministic Case

We first consider the deterministic case where we are given impulse response data or the Markov parameters. Writing the input/output relationship in matrix notation (see sec. III) yields

$$\begin{bmatrix} T_n & 0 \\ H_T & T_T \end{bmatrix} \begin{bmatrix} \mathbf{a}_n \\ 0_T \end{bmatrix} = \begin{bmatrix} \mathbf{b}_n \\ 0_T \end{bmatrix} , \tag{A4}$$

where T_n is a lower-triangular and H_T is a full Töplitz matrix of the Markov parameters (H_T is the column reverse ordered Hankel matrix). Letting $T \to \infty$, we get the normal equation

$$\begin{bmatrix} T_n^T & H_T^T \\ I & 0 \end{bmatrix} \begin{bmatrix} T_n & I \\ H_T & 0 \end{bmatrix} \begin{bmatrix} a & 0 \\ -b_1 & e_1 \end{bmatrix} =$$

$$= \begin{bmatrix} R_n & T_n^T \\ T_n & I \end{bmatrix} \begin{bmatrix} a & 0 \\ -b_1 & e_1 \end{bmatrix} = \begin{bmatrix} e_1 b_0 b_0^T & e_1 b_0 \\ e_1 b_0 & e_1 \end{bmatrix} . \qquad \text{(A5)}$$

Stochastic Case

From a stochastic process point of view we can express the normal equation associated with the augmented AR model as

$$E\left\{ \begin{bmatrix} y_t \\ u_t \end{bmatrix} [\, y_t^T \; u_t^T \,] \begin{bmatrix} a & 0 \\ -b_1 & e_1 \end{bmatrix} \right\} = E\left\{ \begin{bmatrix} y_t \\ u_t \end{bmatrix} [\, u_t b_0 \; u_t \,] \right\}$$

$$= \begin{bmatrix} R_n & T_n^T \\ T_n & I_n \end{bmatrix} \begin{bmatrix} a & 0 \\ -b_1 & e_1 \end{bmatrix} = \begin{bmatrix} e_1 b_0 b_0^T & e_1 b_0 \\ e_1 b_0 & e_1 \end{bmatrix} . \qquad \text{(A6)}$$

We can solve for the normal equation of a_n :

$$R_n a_n = [\, R_n - T_n^T T_n \,] a_n = [\, H_\infty^T H_\infty \,] a_n = e_1 R^\epsilon{}_n . \qquad \text{(A7)}$$

The equations (A5), (A6), and therefore the non-Töplitz equations (A7) (!) can be solved recursively with the LWR algorithm. Note that if $R^\epsilon{}_k = 0$, the minimal order $n = k$. We could bring equations (A5), (A6) into a more familiar form by the interleaving permutation $(1, 3, 5, \ldots, 2n-1, 2, 4, 6, \ldots, 2n+2)$, cf. [MDHV], to convert the two-process covariance matrix into a n by n block Töplitz matrix, with 2 by 2 blocks, however the LWR algorithm clearly applies to both representations with suitable modifications.

Thus we have shown that the joint impulse response/covariance matching

problem is equivalent to solving a set of normal equations associated with a two channel AR modeling problem. Since the predictor for the joint process is *triangular* and *minimum phase*, the denominator a_n of the underlying ARMA model is also *minimum phase* and therefore *stable*, (for all k).

Equations (A5) and the elegant stability proof were actually first obtained by Claerbout [Cla1] via a least-squares rational approximation. The connections between the joint innovations representation, the augmented normal equations, and the Hankel matrix were pointed out in [Mo] and also in [MDKV], [MKD], [DKM], where algorithms were given to solve equations of the type seen in (A6) and (A7). For the special case where R has a "shift-low-rank" of one of the type (B5), called the post-windowing method in [MDKV], a Levinson-type algorithm was given recently by [MR]. The stability property of the AR model was proved there using a somewhat more complicated Lyapunov technique.

Appendix B: LS-Recursions for Ladder Forms

The Prewindowing Case

Given a series of observations $\{y(t), 0 \le t \le T\}$, where $\{ y(\cdot) \}$ can be m vectors, we wish to find the least-squares one-step predictor of order p parametrized by the (matrix) coefficients $\{A_{p,T}(i), i=1, \ldots, p\}$. We can define many different squared error criteria $E_{p,T}$, for instance as a function of s and f in

$$E_{p,T} \;\hat{=}\; \sum_{t=s}^{f} \epsilon^{T}_{p,T}(t)\, \epsilon_{p,T}(t) \quad , \quad \epsilon_{p,t} \;\hat{=}\; A^{T}_{p,T}\, \boldsymbol{y}[t:t-p] \; ,$$

$$A^{T}_{p,T} \;\hat{=}\; [\, \mathbf{I} \, , A^{T}_{p,T}(1), \ldots , A^{T}_{p,T}(p)] \; , \quad \boldsymbol{y}^{T}[t:t-p] \;\hat{=}\; [y_t^{T} , \ldots , y_{t-p}^{T}] \qquad \text{(B1)}$$

An obvious choice from an *innovations* point-of-view is $(s=0, f=T)$, the "pre-windowing" case [MDKV]. If $s = p$ and $f = T$ the so-called "covariance" method is obtained, and if $s = 0$ and $f = T + p$ we get the pre- and post-windowed case or the "correlation" method [MG]. The total squared error can be expressed as

$$E_{p,T} \;=\; tr\, \{\, A^{T}_{p,T}\, \boldsymbol{R}_{p,T}\, A_{p,T} \,\} \quad , \quad \boldsymbol{R}_{p,T} \;=\; \boldsymbol{Y}_{p,T}\, \boldsymbol{Y}^{T}_{p,T} \; ,$$

$$\boldsymbol{Y}_{p,T} \;\hat{=}\; [\, \boldsymbol{y}[0:-p] \, , \, \boldsymbol{y}[1:-p+1] \, , \, \ldots \, , \, \boldsymbol{y}[T:T-p] \,] \qquad \text{(B2)}$$

Thus the problem of determining $A_{p,T}$ by minimizing $E_{p,T}$ leads to

$$\boldsymbol{R}_{p,T}\, A_{p,T} \;=\; [\boldsymbol{R}^{\epsilon}_{p,T} \, , \, 0, \, \ldots \, , 0]^{T} \, , \quad tr\, \boldsymbol{R}^{\epsilon}_{p,T} \;=\; min\; E_{p,T} \; . \qquad \text{(B3)}$$

Although $\boldsymbol{R}_{p,T}$ is not Töplitz, it still carries a certain shift-invariance structure, given by the following identities

$$\begin{aligned} \boldsymbol{R}_{p,T} \;&=\; \boldsymbol{R}_{p,T-1} \;+\; \boldsymbol{y}[T:T-p]\; \boldsymbol{y}[T:T-p]^{T} \\ &=\; \begin{bmatrix} x & x & x \\ x & \boldsymbol{R}_{p-1,T-1} & \end{bmatrix} \;=\; \begin{bmatrix} \boldsymbol{R}_{p-1,T} & x \\ & x \;\; x \;\; x \end{bmatrix} . \end{aligned} \qquad \text{(B5)}$$

Define the backward predictor $B_{p,T}$ and the smoothing errors $C_{p,T}$

$$B^{T}_{p,T}\, \boldsymbol{R}_{p,T} \;\hat{=}\; [\, 0 \, , \ldots , \, 0 \, , R^{r}_{p,T}] \, ; \quad C^{T}_{p,T}\, \boldsymbol{R}_{p,T} \;\hat{=}\; \boldsymbol{y}[T:T-p] \; . \qquad \text{(B6)}$$

Then the forward and backward prediction errors (innovations), $\epsilon_{p,T}$, and $r_{p,T}$, and an auxiliary scalar $\gamma_{p,T}$ can be defined by

$$[\, \epsilon^{T}_{p,T} \, , \, r^{T}_{p,T} \, , \, \gamma_{p,T} \,] \;\hat{=}\; \boldsymbol{y}^{T}[T:T-p]\; [\, A_{p,T} \, , \, B_{p,T} \, , \, C_{p,T} \,] \; .$$

Order Update Recursions

Using the three shift-invariance identities for $\boldsymbol{R}_{p,T}$ (B5) and using some symmetry properties, the order update recursions for $A_{p,T}$, $B_{p,T}$, $C_{p,T}$, $\boldsymbol{R}^{\epsilon}{}_{p,T}$, and $\boldsymbol{R}^{r}{}_{p,T}$ are

$$
\begin{aligned}
A^{T}{}_{p+1,T} &= [\, A_{p,T},\ 0\,]^{T} \;-\; K^{T}{}_{p,T}\, \boldsymbol{R}^{-r}{}_{p,T-1}\, [\, 0,\ B^{T}{}_{p,T-1}\,]^{T} \\
B^{T}{}_{p+1,T} &= [\, 0,\ B_{p,T-1}\,]^{T} \;-\; K_{p,T}\, \boldsymbol{R}^{-\epsilon}{}_{p,T}\, [\, A^{T}{}_{p,T},\ 0\,]^{T} \qquad \text{(B7)} \\
C^{T}{}_{p+1,T} &= [\, C^{T}{}_{p,T},\ 0\,]^{T} \;+\; r^{T}{}_{p+1,T}\, \boldsymbol{R}^{-r}{}_{p+1,T}\, B^{T}{}_{p+1,T} \quad \text{where} \\
K_{p,T} &= [\, \textit{last block row of } \boldsymbol{R}_{p+1,T}\,]\, [\, A^{T}{}_{p,T},\ 0\,]^{T} \\
&= [\, 0,\ B^{T}{}_{p,T-1}\,]\, [\, \textit{first block row of } \boldsymbol{R}_{p+1,T}\,]^{T}. \\
\boldsymbol{R}^{\epsilon}{}_{p+1,T} &= \boldsymbol{R}^{\epsilon}{}_{p,T} \;-\; K^{T}{}_{p,T}\, \boldsymbol{R}^{-r}{}_{p,T-1}\, K_{p,T} \\
\boldsymbol{R}^{r}{}_{p+1,T} &= \boldsymbol{R}^{r}{}_{p,T-1} \;-\; K_{p,T}\, \boldsymbol{R}^{-\epsilon}{}_{p,T}\, K^{T}{}_{p,T}\,. \qquad \text{(B8)}
\end{aligned}
$$

The order update recursions are very similar to the multivariate version of the Levinson algorithm, and a similar set of recursions for time-update can also be obtained [MDKV],[Mo].

Ladder Type Realization

Premultiplying the above equations by $\boldsymbol{y}[T:T-p+1]$, we obtain the following order update recursions for $\epsilon_{p,T}$, $r_{p,T}$, $\gamma_{p,T}$

$$
\begin{aligned}
\epsilon_{p+1,T} &= \epsilon_{p,T} \;-\; K^{T}{}_{p,T}\, \boldsymbol{R}^{-r}{}_{p,T-1}\, r_{p,T-1} \\
r_{p+1,T} &= r_{p,T-1} \;-\; K_{p,T}\, \boldsymbol{R}^{-\epsilon}{}_{p,T}\, \epsilon_{p,T} \\
\gamma_{p+1,T} &= \gamma_{p,T} \;+\; r^{T}{}_{p+1,T}\, \boldsymbol{R}^{-r}{}_{p+1,T}\, r_{p+1,T}\,. \qquad \text{(B9)}
\end{aligned}
$$

The "*Kalman gain*" $K_{p,T}$ is obtained from [MV] as follows

$$
K_{p,T+1} = K_{p,T} + r_{p,T}\, \epsilon^{T}{}_{p,T+1} \;/\; (\,1 - \gamma_{p-1,T}\,) \qquad \text{(B10)}
$$

and the *reflection* or *PARCOR coefficients* are obtained by

$$
K^{\epsilon}{}_{i,T} \;\hat{=}\; K_{i,T}\, \boldsymbol{R}^{-\epsilon}{}_{i,T}\,; \quad K^{r}{}_{i,T} \;\hat{=}\; K^{T}{}_{i,T}\, \boldsymbol{R}^{-r}{}_{i,T-1}\,. \qquad \text{(B11)}
$$

The initial conditions are given by

$$\epsilon_{0,T} = r_{0,T} = y_T \; ; \qquad \gamma_{-1,T} = 0 \; ;$$

$$R^{\epsilon}_{0,T} = R^{r}_{0,T} = \sum_{t=0}^{T} y_t y_t^{T} = R^{\epsilon}_{0,T-1} + y_T y^{T}_T \; ;$$

for $p \geq T$:

$$\epsilon_{p,T} = \epsilon_{T,T} \; ; \quad r_{p,T} = r_{T,T} \; ; \quad \gamma_{p,T} = \gamma_{T,T} \; ;$$

$$R^{\epsilon}_{p,T} = R^{\epsilon}_{T,T} \; ; \; R^{r}_{p,T} = R^{r}_{T,T} \; ; \; K_{p,T} = 0 \; ;$$

$$K_{p,p+1} = y_0 \; \epsilon_{p,p+1}{}^{T} .$$

As the dual to the stochastic forms in [IS], [Wak], [Mo], [SKM], equations (B8)-(B11) are a complete set of *order and time update recursions* to obtain the *exact least-squares* ladder form predictor, which is shown in Figure 1.

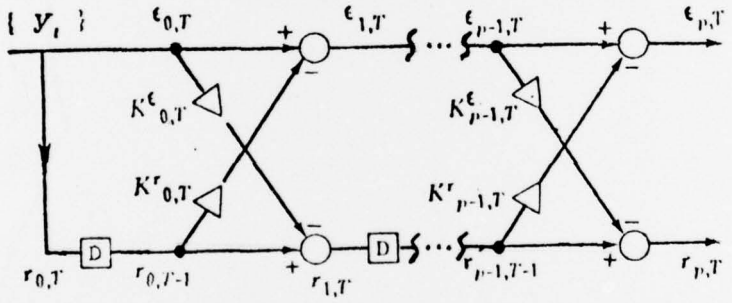

Figure 1. Ladder realization of exact one-step least-squares predictor.

The recursion (B10) computes the *sample cross-covariance* of the *forward and backward innovations*, using the *optimal weighting* $1/(1-\gamma_{\cdot,\cdot})$, compared to other suboptimal schemes [SV]. In the scalar case $R_{p,T}>0$ if $y_0=0$, or in general if $\gamma_{p-1,T}<1$, since $0 \leq \gamma_{p,T} \leq 1$ [MDKV]. If $m>1$, we require $T \geq p+m$. These singularities can be avoided by including a priori estimates of the covariance R_n, or equivalently including a weighted norm of the predictor a_n in the error criteria $E_{p,T}$. Several such modifications have been proven useful in actual implementations.

References

[AE] Astrom,K.J. and P.Eykhoff,"System Identification - A Survey", Automatica,v.7,1971,pp.123-162.

[AH] Atal,B.S.,S.L.Hanauer,"Speech Analysis and Synthesis by Linear Prediction of the Speech Wave," J. Acoust. Soc. Amer., v.5,1971,pp.637-655.

[Aka] Akaike,H.,"Power Spectrum Estimation through Autoregressive Model Fitting," Ann. Inst. Stat. Math., vol. 21,1969,pp.407-419.

[Am] Ambartsumian, V. A.,"Diffuse Reflection of Light by a Foggy Medium," Dokl. Akad. Sci. SSSR, vol.38, 1943, pp.229-322.

[Be] Berlekamp, E.R., *Algebraic Coding Theory*, McGraw - Hill, New York, 1968.

[BJ] Box, G.E., and G. M. Jenkins, *Time Series Analysis Forecasting Forecasting and Control*. San Francisco, Ca.:Holden-Day, 1970.

[Chan] Chandrasekhar, S., Radiative Transfer, Oxford Univ. Press: Oxford 1950; also Dover: New York 1960.

[Cla1] Claerbout, J. F.,"Estimation of a Rational Function on the Unit Circle," unpublished memo.

[Cla2] Claerbout, J. F., *Fundamentals of Geophysical Data Processing*, McGraw-Hill: New York, 1976.

[Di] Dickinson, B. W.,"Two Recursive Estimates of Autoregressive Models Based on Maximum Likelihood," submitted for publication.

[DKM] Dickinson, B., T. Kailath and M. Morf, "Canonical Matrix Fraction and State-Space Descriptions for Deterministic and Stochastic Linear Systems," IEEE Transactions of Automatic Control, vol. AC-19, 1974, pp. 656-667.

[DMK] Dickinson, B., M. Morf and T. Kailath, "A Minimal Realization Algorithm for Matrix Sequences," IEEE Transactions on Automatic Control, vol. AC-19, 1974, pp. 31-38.

[FMKL] Friedlander, B., M. Morf, T. Kailath and L. Ljung, "New Inversion Formulas for Matrices Classified in Terms of Their Distance from Toeplitz

Matrices,"submitted to SIAM J. A. Math.

[GGMS] Gill, P.E., G.H. Golub, W. Murray, and M.A. Saunders, "Methods for Modifying Matrix Factorizations," Stanford Univ. Computer Science Rept. STAN-CS-72-322, 1972.

[GK] Gevers, M., and T. Kailath, "An Innovations Approach to Least - Squares Estimation, Part VI," IEEE Trans. Automat. Contr. vol. AC-18, 1973, pp. 588-600.

[GrMo] Gray, R. M.,"Source Encoding," lecture by M. Morf, ISL Report, NSF-Workshop, Stanford University, Summer 1974.

[GS] Grenander, U., and G. Szegö, *Toeplitz Forms and Their Applications*. Berkeley, Calif.:Univ. California Press,1958.

[HK] Ho,B.L., and R.E. Kalman, "Effective Construction of Linear State - Variable Models from Input - Output Functions," Regelungstechnik, vol. 14, 1966, pp. 545-548.

[Hou] Householder, A. S., *The Theory of Matrices in Numerical Analysis*. Blaisdell, New York: 1964.

[IS] Itakura,F., and S. Saito,"Digital Filtering Techniques for Speech Analysis and Synthesis," Conf. Rec., 7th Int. Congr. Acoust.,Budapest,1971,Paper 25 C 1.

[KFA] Kalman, R.E., P.L. Falb and M.A. Arbib,*Topics in Mathematical System Theory*, McGraw-Hill, New York: 1969.

[KKM] Kung. S, T. Kailath, and M. Morf, "A Fast Projection Method for Canonical Minimal Realization," Proc. IEEE Conf. on Decision and: Control, pp.1301-1302, Dec. 1976.

[KVM] Kailath, T., A. Vieira and M. Morf, "Inverses of Toeplitz Operatiors, Innovations, and Orthogonal Polynomials," submitted for publication.

[Lanc] Lanczos, C.,"An Iteration Method for the Solution of the Eigenvalue Problem of Linear Differential and Integral Operators," J.Res.Nat. Bur.Standards, vol.45,1950,pp.255-282.

[Lev] Levinson, N.,"The Wiener RMS (root mean square) Error Criterion in

Filter Design and Prediction," J. Math. Phys.,vol.25,1947, pp.261-278.

[LKF] L.Ljung,T.Kailath,B.Friedlander,Proc. IEEE,v.63,1976,pp.131-139.

[LKM] Lévy,B.,S.Y.Kung,M.Morf,Multi.Dim.Sys.Conf.,Monterey,Dec.1976.

[Mass] Massey, J.L., "Shift-Register Synthesis and BCH Decoding," IEEE Trans. on Inform. Theory, vol. IT-15, 1969, pp. 122-127.

[Makh] Makhoul, J. "Linear Prediction: A Tutorial Review," Proc. IEEE, vol. 63, 1975, pp. 561-580.

[MDKV] Morf, M., B. Dickenson, T. Kailath and A. Vieira, "Efficient Solution of Covariance Equations for Linear Prediction," submitted to IEEE-ASSP 1976.

[MG] Markel, J.D., and A.H. Gray, Jr.,*Linear Prediction of Speech*, Springer-Verlag, Berlin: 1976.

[MK] Morf, M., and T. Kailath, "Square-Root Algorithms for Least- Squares Estimation ,"IEEE Trans.on Automat. Control,vol. AC-20, 1975, pp. 487-497.

[MKD] Morf, M., T. Kailath, and B. Dickinson, "General Speech Models and Linear Estimation Theory," in Speech Recognition, pp.157-182, New York:Academic Press, 1975.

[MKL] Morf, M., T. Kailath and L. Ljung, "Fast Algorithms for Recursive Identification," Proc. IEEE Conf. on Decision and Control,pp.916-921,Dec. 1976.

[MKLNV] Morf, M., T. Kailath, D.T. Lee, J.R. Nickolls, and A. Vieira, "A Classification of Algorithms for ARMA Models and Ladder Realizations", Proc. IEEE Conf on Acoustics, Speech, and Signal Processing, Hartford, CT, May, 1977.

[Mo] Morf, M., "Fast Algorithms for Multivariable Systems," Ph.D. dissertation, Stanford University, Stanford, California 1974.

[MR] Mullis, C. T., and R. A. Roberts, "The use of Second-Order Information in the Approximation of Discrete-Time Linear Systems," IEEE Trans. Acoust. Signal Processing, vol. ASSP-24,1976, pp.226-238.

[MV1] Morf. M., A. Vieira "Recursive Least-Squares Estimation of Partial Correlations," submitted to IEEE-ASSP 1977.

[MV2] Morf. M., A. Vieira "Multichannel Least-Squares Estimation of Partial Correlations," submitted to IEEE-AC, 1977.

[MVK] Morf. M., A. Vieira and T. Kailath, "The Multi-channel Maximum Entropy Method," submitted for publication.

[Par] Parzen, E.,"Multiple Time Series Modeling," in *Multivariate Analysis 2*, ed. by P. R. Krishnaiah, pp.389-409, Academic: New York, 1969.

[Rob] Robinson, E. A.,*Multichannel Time Series Analysis with Digital Computer Programs* San Francisco:Holden-Day, 1967.

[SKM] Sidhu, G. S., T. Kailath and M. Morf, "Development of Fast Algorithms via Innovations Decompositions," Proc. of 7th Hawii Intl. Conf. on Inf. and System Sciences, Honolulu, Hawaii, pp 192-195, Jan. 1974.

[Si] Silverman, L., "Realization of Linear Dynamical Systems," IEEE Trans. Automat. Contr., vol. AC-16, 1971, pp. 554-567.

[SLG] Soderstrom, T., L.Ljung, and I. Gustavsson, "A Comparative Study of Recursive Identification," Rep. 7427, Div. of Auto. Contr., Lund Inst. of Tech., 1974.

[SV] Srinath, M.S.,M. M. Viswanathan,"Sequential Algorithm for Identification of Parameters of an Autoregressive Process," IEEE Trans. Aut. Contr.,AC-20,1975,pp.542-546.

[Sze] Szegö,G.,"Orthogonal polynomials," Amer. Math. Soc. Colloq. Publ., vol. 23,1939; 2nd ed., 1958; 3rd ed. 1967.

[UB] Ulrych, T. J., T. N. Bisshop,"Maximum Entropy Spectral Analysis and Autoregressive Decomposition,"Rev. Geophys. Space Phys.,v.13, 1975,pp.183-200.

[Wak] Wakita, H., "Estimation of the Vocal Tract Shape by Optimal Inverse Filtering and Acoustic/Articulatory Conversion Methods," Monograph No.9, Speech Communications Res. Lab. Inc., Santa Barbara, Calif., July 1972.

[Wie] Wiener N., *Time Series*,M.I.T. Press, Massachusetts Institute of Technology, Cambridge, Massachusetts, 1949.

[WR] Wiggins, R.A.and E.A. Robinson,"Recursive Solution to the Multichannel Filtering Problem," J. Geophys. Res., vol. 70,1965, pp.1885-1891.

[YT] Youla, D.C. and P. Tissi, "N-Port Synthesis via Reactance Extraction Part I," IEEE Int. Convention Record, vol. 14, pt. 7, 1966, pp. 183-205.

APPENDIX II

PROJECT PUBLICATIONS

PROJECT PUBLICATIONS

[1] "Application of Frequency Discrimination Technique to the Analysis of Electroencephalographic Signals", D.C. Lai and R.L. Lux, Proc. National Electronics Conference, vol. 27, 80-85, October 1972.

[2] "A Model for the Photically Stimulated Electroencephalographic Signals", R.V. Floyd, J.E. Anliker, and D.C. Lai, Proc. 1973 San Diego Biomedical Symposium, February 1973.

[3] "The Graphics Software DEC Forgot to Include", K.H. Jacker, Proc. of the DECUS Spring Symposium, May 2, 1973.

[4] "Estimating Signal and Noise in Coherent Time Averages of EEG Data", J.E. Anliker, D.C. Lai, T. Rimmer, and H. Finger, Proc. 1973 Annual Conference on Engineering in Medicine and Biology, October 1973.

[5] "Real-time EEG Analysis and Monitoring Using In-phase and Quadrature Components", M. Ein-Gal and D.C. Lai, Proc. 26th ACEMB, p.401, October 1973.

[6] "Error-free EEG Signal Representation", M. Ein-Gal and D.C. Lai, Proc. IEEE International Conference on Systems, Man, and Cybernetics, pp.242-243, November 1973.

[7] "Computer Determination of Eye Fixations and Saccades", A. Shah and D.C. Lai, Proc. 27th ACEMB, p.103, October 1974.

[8] "A Nonlinear Model of EEG Entrainment by Periodic Photic Stimulation", J.R. Nickolls, D.C. Lai, and J.E. Anliker, Proc. 7th Annual Meeting of the Neuroelectric Society, pp.13-14, November 1974.

[9] "Remark on Algorithm 479[Z]", H.S. Magnuski, Communications of the ACM, vol. 18, no. 2, p.119, February 1975.

[10] "Prediction of EEG Alpha Waveform by Using an Autoregressive Model", A. Shah, D.C. Lai, and J.E. Anliker, Proc. 1975 San Diego Biomedical Symposium, February 1975.

[11] "Computer Control of the Foveal and Peripheral Visual Field", H.S. Magnuski and D.C. Lai, Proc. 28th ACEMB, September 1975.

[12] "Modeling the EEG Entrainment Process", J.R. Nickolls, D.C.Lai, and J.E. Anliker, Proc. 28th ACEMB, September 1975.

[13] "Monitoring and Prediction of Saccadic Eye Movement", A. Shah and D.C. Lai, Proc. 1975 IEEE International Conference on Systems, Man and Cybernetics, September 1975.

[14] "Modeling of Saccadic Eye Movements and EEG Alpha Rhythm", Arun Shah, Ph.D. dissertation, Department of Electrical Engineering, Stanford University, Stanford, California, 1977.

[15] "A Nonlinear Model of the Human EEG Signal during Photic Stimulation", J.R. Nickolls, Ph.D. dissertation, Department of Electrical Engineering, Stanford University, Stanford, California, 1977.

[16] "Alpha Phase Contingency Analysis of the Average Visual Evoked Response:, James E. Anliker, submitted for publication in Electroencephalography and Clinical Neurophysiology.

[17] "Alpha Phase Probability Analysis of the Average Visual Evoked Potential", J.E. Anliker and R.V. Floyd, San Diego Biomedical Symposium, 1977, (in press).

[18] "Eye Movements: On-line Measurement, Analysis, and Control", J.E. Anliker, A chapter (pp.185-202) in Monty, R.A. and J. Senders (Eds.), Eye Movements and Psychological Processes. Hillsdale: Lawrence Erlbaum Associates, 1976.

[19] "Biofeedback from the Perspectives of Cybernetics and Systems Science", J.E. Anliker, A chapter in Beatty, J. and H. Legewie (Eds.), Biofeedback and Behavior: A NATO Symposium. N.Y.: Plenum, 1977 (in press).

[20] "A Classification of Recursive Modeling Methods", M. Morf and T. Kailath, submitted for publication in IEEE Transactions on Acoustics, Speech & Signal Processing.